Labor and Public Studies

Labor and Public Studies

A dissertation presented

by

Jessica Laird

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2017

Dissertation Advisors:
Professor Larry Katz
Professor Nathan Hendren

Author:
Jessica Laird

Labor and Public Studies

Abstract

The first chapter studies how firms set contributions to employer-provided 401(k)-type pension plans. Using a reform that decreased the subsidy for contributions to capital pension accounts for Danish workers in the top income tax bracket, we provide strong evidence that employers' contributions are based on their employees' savings preferences. We find an immediate decrease in employer contributions to capital accounts, whose magnitude increased in the share of employees directly affected by the reform.

The second chapter studies the impacts of physicians with differential prescribing behaviors on patient prescription drug use and labor market outcomes for the four classes of prescription drugs used most frequently to treat musculoskeletal and mental health disorders: opioids, anti-inflammatories, anti-anxiety drugs, and anti-depressants. I use Danish administrative data and exploit quasi-random separations of individuals from their physicians associated with geographic moves across municipalities to estimate the causal impact of physician prescribing rates on individual prescription drug use and labor market outcomes. I find that having a general practitioner who has a 10 percentage point higher opioid prescription rate leads to a 4.5 percentage point increase in the probability an individual uses prescribed opioids, as well as a 1.5 percentage point decrease in their labor force participation. Changes in physician prescribing rates lead to similar changes in prescription drug use for the other classes of prescription drugs, but they are not associated with any discernible effect on labor market outcomes.

In the third chapter, I use Denmark Administrative population data to identify exposure effects of municipalities on mortality. Using variation from individuals who move between 1984-2004, we find that each additional year an individual lives in a municipality in 1984-2004 increases the association between their probability of death from 2005-2015 and the municipality's

10 year mortality rate by 3 percentage points. When we specifically estimate exposure effects, we find that living in a municipality with a 1 SD higher exposure mortality effect for 60 years increases an individual's 10 year mortality rate by 3.4 percentage points. This effect attenuates by approximately 10% when we include controls, but still remains large and significant.

Contents

Chapter 1

Do Employer Pension Contributions Reflect Employee Preferences? Evidence From A Retirement Savings Reform In Denmark[1]

1.1 Introduction

With the decline in the prevalence of defined-benefit pension plans, individual savings in defined-contribution accounts are becoming an increasingly important income source for post-retirement consumption. A large and growing portion of savings balances in defined-contribution accounts is within employer-sponsored pension plans, such as 401(k)s.[2] Recent research has underlined the important role that employers' decisions play in determining employees' actual savings within employer-sponsored accounts, since most workers do not actively deviate from the default options which are set by their employer (Madrian and Shea 2001; Choi et al. 2004; Choi, Laibson, and Madrian 2007; Beshears et al. 2009; Gelber 2011). The impact of employers' decisions has also been

[1] Co-authored with Itzik Fadlon and Torben Nielsen

[2] In 2014, 35 percent ($8.3 trillion) of retirement assets in the US were held in defined-benefit plans, while a much larger share of 58 percent ($13.8 trillion) was held in defined-contribution accounts: 30 percent in Individual Retirement Accounts (IRAs) and 28 percent in employer-sponsored plans, mostly 401(k)s (Choi 2015).

found to translate into large effects on individuals' overall level of savings (Chetty et al. 2014). This research has led policymakers to consider introducing policies that encourage employer contributions to pension accounts in order to increase individuals' retirement savings (Beshears et al. 2010).

But how effectively private firms represent their employees' savings interests remains an open question. In theory, both standard models and models of altruistic planners suggest that firms should have incentives to make efficient savings choices on their workers' behalf. For example, standard models of efficient compensation arrangements by competitive firms (as in Rosen 1974), predict that employers will provide benefits when firms can purchase goods or services more cost-effectively than employees, and will provide the optimal package that their workers will value most highly.[3] Similarly, models of purely paternalistic firms that incorporate individual optimization frictions predict that firms will provide their employees with optimal retirement savings plans, taking into account their employees' suboptimal behavior (Choi et al. 2003; Cremer et al. 2008; Carroll et al. 2009; Cremer and Pestieau 2011; Goda and Manchester 2013; Roeder 2014; Fadlon and Laibson 2016)

However, firms' incentives or ability to represent their employees' savings interests may be weakened for a variety of reasons. First, if employees are inattentive or "unsophisticated" (in the sense that they are not aware of their suboptimal decision making), they may fail to recognize the value of firms' choices on their behalf. This is a likely possibility given the evidence that most employees are inattentive with regards to their own savings and may have only imperfect knowledge about their employer pension plans (Mitchell 1988; Gustman, Steinmeier, and Tabatabai 2009). Second, firms' incentives may be additionally offset by the costs of efficiently managing their employees' pensions. These costs could be substantial, since managing pensions requires monitoring policy changes that alter tax incentives for contributions to retirement accounts, as well as keeping track of how these changes differentially affect the firms' heterogeneous workforce. Finally, potential mismanagement and inattention within the firm may lead to inefficient savings

[3]In our application, provision of pension benefits is likely less costly for the employer since there are often economies of scale with respect to the acquisition costs and managing of the pension product (e.g., in terms of average fees to the financial service provider per saver – see Danish Ministries of Business and Economics, Finance, Employment and Taxation 2003). Additionally, there is evidence that it is costly for individual households to optimally choose complex financial products due to low levels of financial literacy (Brobeck 1990; Bernheim 1998; Hilgert, Hogarth, and Beverly 2003; Lusardi and Mitchell 2007; Hastings, Madrian, and Skimmyhorn 2013).

choices by firms themselves. Put together, these different factors imply that it is unclear whether the costs of optimally designing retirement savings plans on the workers' behalf outweigh the incentives for firms to do so.

Beyond the question of *whether* firms tailor pension plans to employees' saving preferences, it is important for policy design to study the nature of firm responses to changes in these preferences. The increased reliance on employer-based savings as a source of consumption after retirement necessitates understanding empirically *how fast* and to *what extent* employers respond to shocks to the economic environment that can alter employees' savings incentives.

In this paper, we empirically analyze how firms set characteristics of their contributions to employer-sponsored pension plans in practice.[4] To do so, we exploit a reform to the Danish retirement savings system. This reform differentially affected employees according to their exact location on the labor-income tax schedule and differentially changed tax deductions for contributions to two types of savings accounts: "capital" savings accounts, which are paid out in full at retirement, and "annuity" savings accounts, which are paid out as an annuity. Specifically, in 1999, the Danish government decreased the subsidy to contributions to capital pension accounts for workers in the top income tax bracket, while the subsidies to capital pension contributions for workers in lower tax brackets and to annuity pension contributions remained unchanged.

We find that immediately following the reform, employers significantly decreased their annual contributions to capital pension accounts. The average decrease was on the order of 27 percent – 0.76 percentage points (pp) on a baseline contribution rate of 2.81 pp. This decrease was entirely driven by firms in which some share of the workforce was directly affected by the reform, with no responses in workplaces in which all employees had earnings below the top income tax bracket. Moreover, the response strongly and continuously increased in the share of employees above the top tax threshold. We find that an additional 10 percent of workers at the top bracket led to an additional decrease of more than 0.2 pp in employer contributions, so that workplaces in which all employees were at the top bracket experienced a significant drop of more than 2 pp on a base

[4]To our knowledge, this is the first paper that studies this specific topic. Some papers analyze firms' choice of whether to offer a pension plan (or the choice between defined-benefit or defined-contribution plans), but these papers usually focus on supply-side factors that affect firms' cost of providing the plan (Aaronson and Coronado 2005; Dummann 2008; Hernaes et al. 2011). Papers that do analyze demand-side factors, relate these to the individual take-up of plans rather than to the firms' decision to offer them (Aaronson and Coronado 2005; Dummann 2008).

of 3.5 pp.

In order to put the employers' responses and their magnitude in context, the Danish setting allows us to compare contributions within employer-provided 401(k)-like accounts to individual contributions within private IRA-like accounts, which were equally affected by the reform. We find that for individuals at the top bracket almost the entire change in overall capital retirement savings was attributable to employer responses. We also show that the clear gradient of changes in employer-provided accounts with respect to the share of directly affected individuals in the workplace disappears in the analysis of changes in private accounts, suggesting that employer responses were not crowded out by individual responses in other closely substitutable accounts.

By changing the relative prices of contributions to capital and annuity accounts, the reform rendered contributions to annuity accounts more financially attractive through a substitution effect, but also led to an income effect that would push toward an overall decrease in pension contributions. Studying employer contributions to annuity accounts, we find that employers compensated for the decrease in capital contributions with an equally-large increase in annuity contributions, with no decrease in total pension contributions.[5] This suggests that at the employer-level the effect was driven by a substitution effect. In fact, annuity accounts serving as a close substitute is likely the reason there was such a large response in capital accounts. We also show that the potential income effect of the reform had no effect on other means of employee compensation, namely, labor income.

Lastly, we provide a suggestive analysis for assessing the optimality of firms' responses to the reform from the perspective of the employees. As a benchmark for "optimal" responses, we use the actions taken by "attentive" individuals, who made changes to their self-managed IRA-like savings accounts in response to the reform. The analysis supports the notion that larger firms make more optimal decisions compared to smaller firms, and that delegating decisions to employers may be welfare increasing.

The paper proceeds as follows. In Section 1.2, we discuss the institutional setting of the policy change and the data that we use. In Section 1.3, we provide the empirical analysis of employers' responses to the reform and their heterogeneity with respect to workplace composition.

[5]These employer responses are in contrast to individual responses in private accounts, for which Chetty et al. (2014) find a shift of 57 cents to annuity accounts for each DKr that individuals would have contributed to capital accounts.

Specifically, Section 1.3.1 analyzes employer responses in contributions to capital pension accounts and their timing; Section 1.3.2 puts the magnitude of these responses in context by comparing the changes in contributions to employer-provided accounts to individuals' responses within their private accounts; Section 1.3.3 studies employers' shift to contributing to annuity accounts; and Section 1.3.4 provides a simple analysis of the optimality of the employers' responses to the reform. Section 1.4 concludes.

1.2 Institutional Details and Summary of Data

1.2.1 Institutions

This section provides the necessary background on Danish retirement institutions that is important for our empirical analysis.[6] In Denmark, there are two types of defined-contribution (DC) pension savings accounts similar to the US – employer-sponsored accounts, similar to 401(k)s, and private accounts, similar to IRAs. Employer-sponsored and private DC accounts have equivalent tax properties but are completely independent, which makes them close substitutes. Within both the employer-sponsored and the private DC pension plans, there are two types of tax-preferred accounts: capital pension accounts and annuity pension accounts. Capital pension accounts are paid out as a lump sum and taxed at 40 percent on payout, while annuity pension accounts are paid out over several years and are taxed as personal income. Balances in capital pension accounts can be converted to annuity pensions when they become liquid, but the reverse is not allowed. Contributions to both types of accounts are tax deductible at the time of contribution (as in traditional non-Roth 401(k)s and IRAs), and capital gains are taxed at 15 percent, compared to approximately 29 percent for assets in taxable accounts.

Our empirical research design exploits a 1999 tax reform, within which the Danish government aimed at reducing the generosity of capital accounts and incentivizing a shift to annuity accounts. To do so, the reform reduced the average deduction for contributions to capital pensions from 59 cents per DKr to 45 cents per DKr for individuals in the top income tax bracket. The deduction

[6]For additional information see OECD (2009) or Bingley, Gupta, and Pedersen (2007).

for those in the lower tax bracket remained the same at 45 cents per DKr.

There were additional policy changes associated with this reform, some of which affected individuals differentially across the income distribution.[7] However, most importantly for our purposes, these additional changes did not affect the interaction between individual income and the tax treatment of contributions to capital savings accounts or annuity accounts. Since our research design relies on this interaction, we are confident that the results reflect responses to the reduction in the capital contributions subsidy for individuals at the top income tax bracket.[8]

Most jobs in Denmark (roughly 80 percent) are covered by collective bargaining agreements between worker unions and employer associations. These agreements often have a pension plan in which a fixed proportion of an individual's earnings is paid into a retirement account. For the 20 percent of jobs that are outside the common agreements, employers set contribution rates to capital and/or annuity accounts for their workers.[9] While individuals cannot change the total contribution rate, they can choose a different allocation across capital and annuity accounts, but only if their pension fund allows both types of accounts.

1.2.2 Data Sources, Sample Selection, and Variable Definitions

We merge data from several administrative registers of the Danish population – the income tax register, the population register, and the Danish Integrated Database for Labor Market Research (IDA) – to obtain annual information on Danish employees and their matched firms from 1996 to

[7]These changes include a reduction in the deduction value of negative capital income, a reduction in the bottom bracket tax rate, a move to equalize taxation on all liquid assets (i.e., stocks vs. bonds), a decrease in the value of the Voluntary Early Retirement Plan, and the possibility to initiate a private rate-pension plan (a special type of an annuity pension plan) after age 55. While the possibility to initiate a private rate pension plan may have affected the incentives for initiating annuity pension plans for older workers, effectively, we find no differential savings patterns around the age threshold of 55. Hence, this component of the reform is not affecting our results. The remaining reform elements may affect people differentially across the income distribution, but do not directly affect the demand for capital and annuity contributions.

[8]Nevertheless, to alleviate concerns regarding the other potential confounding policy changes, in the empirical section we include a fifth order polynomial of individual income as control variables in our regression analysis and find that our results do not change.

[9]The contributions rates, default portfolio allocations, and administration fees are set by bargaining between the pension fund and the employer, which is usually represented by the heads of HR departments, CFOs, and pension brokers. Updates to the employees' plans, e.g., in response to taxation changes, are made by the pension fund and the employer.

2001. These registers include data on taxable labor earnings, contributions to pension accounts, occupation, industry, and employees' demographics (such as age and educational attainment). All income and savings variables used in the analysis are based on third-party reports: earnings and pension contributions are reported directly by employers and pension funds to the tax authority.

Starting from the population data-set, we impose four restrictions to obtain our primary analysis sample. First, we exclude individuals under age 20 or over age 60, at which the majority of the Danish workforce is eligible for early retirement benefits and retirement savings are eligible for withdrawal (without a penalty). Second, we focus our analysis on the 20 percent of workers who are outside collective bargaining agreements for which contribution rates are set by the firm (rather than within collective bargaining agreements).[10] To isolate the jobs that are not covered by collective bargaining, we exclude workers in the public sector or in blue-collar occupations, since they are likely covered by collective agreements.[11] Therefore, our analysis sample consists of workers in private firms with white-collar occupations.[12] Third, we exclude observations of workers with self-employment income because their "employer" contributions are not set by a firm. Finally, we exclude occupation-firm cells with fewer than five employees in order to decrease measurement error, as such small cells are unlikely to be treated as an independent unit by employers.[13] Overall, our sample choice allows us to study how private firms in competitive markets design their employees' pension plans.

We run our analysis at the occupation-firm level, since firms often set contribution rates separately by occupation as the pension funds in Denmark (that the firms negotiate with when setting their retirement plans) are largely occupation based. Specifically, we differentiate occupations at the 2-digit ISCO occupation code level.[14] We measure contribution rates to employer-sponsored

[10]While it would be interesting to additionally analyze how pension plans are designed within a collective bargaining setting, the data do not allow us to match workers to unions and firms to employer associations.

[11]See Appendix A.5 for a complete description of occupations that we define as white- or blue-collar.

[12]Still, some white-collar jobs in the private sector are covered by collective bargaining. Therefore, in Appendix A.4, we assess how inadvertently including workers covered by collective agreements may affect the results and show that it likely only attenuates our estimates.

[13]Our results are not sensitive to this choice – see Appendix A.1 for analyses that vary this minimal cell-size restriction.

[14]Due to measurement error in many-digit occupation codes, our choice for the analysis is the 2-digit code level. However, actual employer contributions may be set at a higher- (or lower-) digit occupation code level. This causes

accounts as contribution levels divided by taxable labor income. This measure of contribution rates may vary within an occupation-firm cell since employers may set pension contributions at a finer level within the firm than the 2-digit code that we use, and since individuals can choose a different distribution of contributions between capital and annuity accounts than the employer's default when employers offer both accounts. Therefore, to identify the default contribution rates chosen by employers, we use the median contribution rate within an occupation-firm cell as our measure. In Appendix A.1, we assess the sensitivity of our results to other measures of defaults (namely, modes) and find very similar results. For ease of discussion, we refer in the remainder of the paper to a 2-digit occupation-firm cell as a "workplace" and to the median contribution rate within a workplace (in a given year) as the "employer contribution rate" (or the "default").

1.2.3 Summary Statistics

Table 1.1 presents summary statistics for the sample of private white-collar wage earners between ages 20-60, in workplaces with at least five employees.[15] Our sample contains 2,020,705 worker-year observations from 1996-2001. These amount to 84,764 workplace-year observations with a total of 26,775 unique workplaces.

To provide an overview of contributions to retirement savings accounts in Denmark prior to the reform, Table 1.1 reports information on individual-level pension contributions for the year 1998.[16]

Before the reform in 1999, contributions to employer-sponsored capital accounts were on average 3.6 percent of annual earnings, where 66 percent of workers had positive contributions to these accounts. The average of contributions to employer-sponsored annuity accounts was similarly at 3.7 percent of labor income, where 77 percent had positive contributions to these accounts. In contrast, individual contributions to both capital and annuity private pension accounts

---

some measurement error in our identified decision unit. Our results stay similar if we aggregate occupations at the 1-digit level, or even at the firm level.

[15] During our sample period, 57 percent of wage earners were in the private sector, and 70 percent had white-collar occupations.

[16] All monetary values are reported in nominal Danish Kroner (DKr), where the exchange rate during this time period was approximately DKr 6.5 per US $1.

Table 1.1: *Summary Statistics of Analysis Sample*

	Mean (1)	Median (2)	SD (3)
Individual-Level Variables:			
Labor Earnings (DKr)	285,740	253,893	214,500
Pension Contributions before the Reform			
Employer-Sponsored Accounts			
Capital Contributions (DKr)	11,665	6,679	12,490
Capital Contribution Rate (percent)	3.6	2.8	3.8
Percent with Capital Contributions	66		
Annuity Contributions (DKr)	12,604	5,047	17,830
Annuity Contribution Rate (percent)	3.7	1.8	4.5
Percent with Annuity Contributions	77		
Private Accounts			
Capital Contributions (DKr)	3,125	0	7,627
Capital Contribution Rate (percent)	1.1	0	2.6
Percent with Capital Contributions	28		
Annuity Contributions (DKr)	1,470	0	5,859
Annuity Contribution Rate (percent)	0.4	0	1.7
Percent with Annuity Contributions	15		
Workplace-Level Variables:			
Percent of Employees above Top Tax Threshold	49	50	35
Number of Employees	22	9	97
Number of Worker-Year Observations		2,020,705	
Number of Workplace-Year Observations		84,764	
Number of Workplaces		26,775	

Notes: This table presents means, medians, and standard deviations of key variables in our analysis sample of white-collar workers in private-sector firms from 1996 to 2001. The classification of white-collar occupations is described in detail in Appendix E. All monetary values are reported in nominal Danish Kroner (DKr), where the exchange rate during this time period was approximately DKr 6.5 per US $1. Labor income is calculated as total pre-tax wage earnings plus employer pension contributions. The values reported in the table for pension contributions before the reform are based on data from 1998. Pension contribution levels are winsorized at their 99th percentile.

were much lower, with average contribution rates of 1.1 percent and 0.4 percent, respectively.

Importantly for our design, 49 percent of workers were above the top labor-income tax threshold, with a sizable standard deviation of the fraction of workers above the threshold across workplaces on the order of 35 percent. Given the restriction to at least five workers per occupation-firm cell, the average workforce size is 22, while the median cell size is 9.

1.3 Empirical Evidence

In this section, we analyze how employer contributions to pension accounts responded to the reform, namely, to the decrease in subsidies to capital pension contributions for workers with labor income in the top tax bracket. We begin by analyzing changes in employer contributions to capital accounts and their sensitivity to the share of workers who were directly affected by the reform. Then, we assess the magnitude of these responses to the reform by comparing employer responses to individual responses within private retirement accounts. We additionally explore other potential margins of firm responses, in particular, whether changes in capital contributions translate into changes in overall savings or whether firms substitute contributions to annuity accounts. Finally, we briefly discuss the optimality of the firms' responses to the reform.

1.3.1 Employer Responses in Contributions to Capital Pension Accounts

In the years preceding the reform, employer capital contributions steadily increased, such that they were on average 2.81 percent in 1998 (see Panel A of Appendix Figure A.1).[17] In contrast, in 1999, when the capital subsidy decreased for workers above the top threshold, the average employer capital contribution rate decreased by 0.76 percentage points (pp). However, this average drop of 27 percent aggregates the responses of all the employers in our sample of private firms and white-collar occupations. Since the reform changed the savings incentives only for employees in the top labor-income tax bracket, our analysis focuses on the heterogeneity of firm responses

[17]Note that these are employer contributions (measured by workplace-level medians) as opposed to individual-level contributions to employer-sponsored accounts that are reported in Table 1.1.

with respect to the share of the workforce that was directly affected by the reform.

Graphical Analysis. To test whether and to what extent employers' capital contribution responses to the reform increased in the share of workers above the top income threshold, we divide workplaces into equal-sized groups by the fraction of employees above the threshold within a workplace. We begin by plotting in Figure 1.1 the mean employer capital contribution rate against the mean fraction of employees above the top threshold for each group in years 1996-2001. Panel A shows that before the reform, employer capital contributions were increasing in the fraction of workers above the threshold and that the slopes of this relationship were similar across years.[18] However, immediately following the reform – that took effect in 1999 – there is a significant change in this relationship, such that employer capital contribution rates became largely decreasing in the fraction of workers above the threshold. The decrease in employer capital contributions after the reform, i.e., the vertical distance between the lines of years 1998 and 1999, is noticeably larger for workplaces with a higher fraction of directly affected workers.

To clearly see these changes, Panel B of Figure 1.1 displays the year-to-year differences in employer capital contributions as a function of the fraction of employees above the threshold for each year from 1996 to 2001. This figure shows that annual changes in contributions were uniform across different shares of workers above the top tax threshold in the years prior to the reform. However, between years 1998 and 1999, the year of the reform, workplaces with no affected workers did not change their contributions, while those with a larger share of affected employees decreased their contribution rates in larger magnitudes. This change continuously increased in the share of the workforce at the top bracket, with about a 1.9 pp decrease for workplaces with the highest share of affected employees. This is a large response, since completely exiting capital accounts in workplaces with the highest share of affected employees – which is an upper bound to their response – would imply a 3.5pp reduction in contributions. In the subsequent years (2000 and 2001), there were some delayed responses to the reform, but the gradient with respect to the share of affected workers in those years is much smaller.

Overall, the graphical evidence clearly reveals that employers with a greater share of affected

[18]This is consistent with the fact that top-bracket workers enjoyed a larger subsidy to capital contributions on the margin, but as it is a cross-sectional relationship, there is a variety of other reasons for this pattern such as different preferences for savings across individuals with different labor income levels.

Figure 1.1: *Employer Contributions to Capital Pension Accounts by the Share of Workers above the Top Tax Threshold*

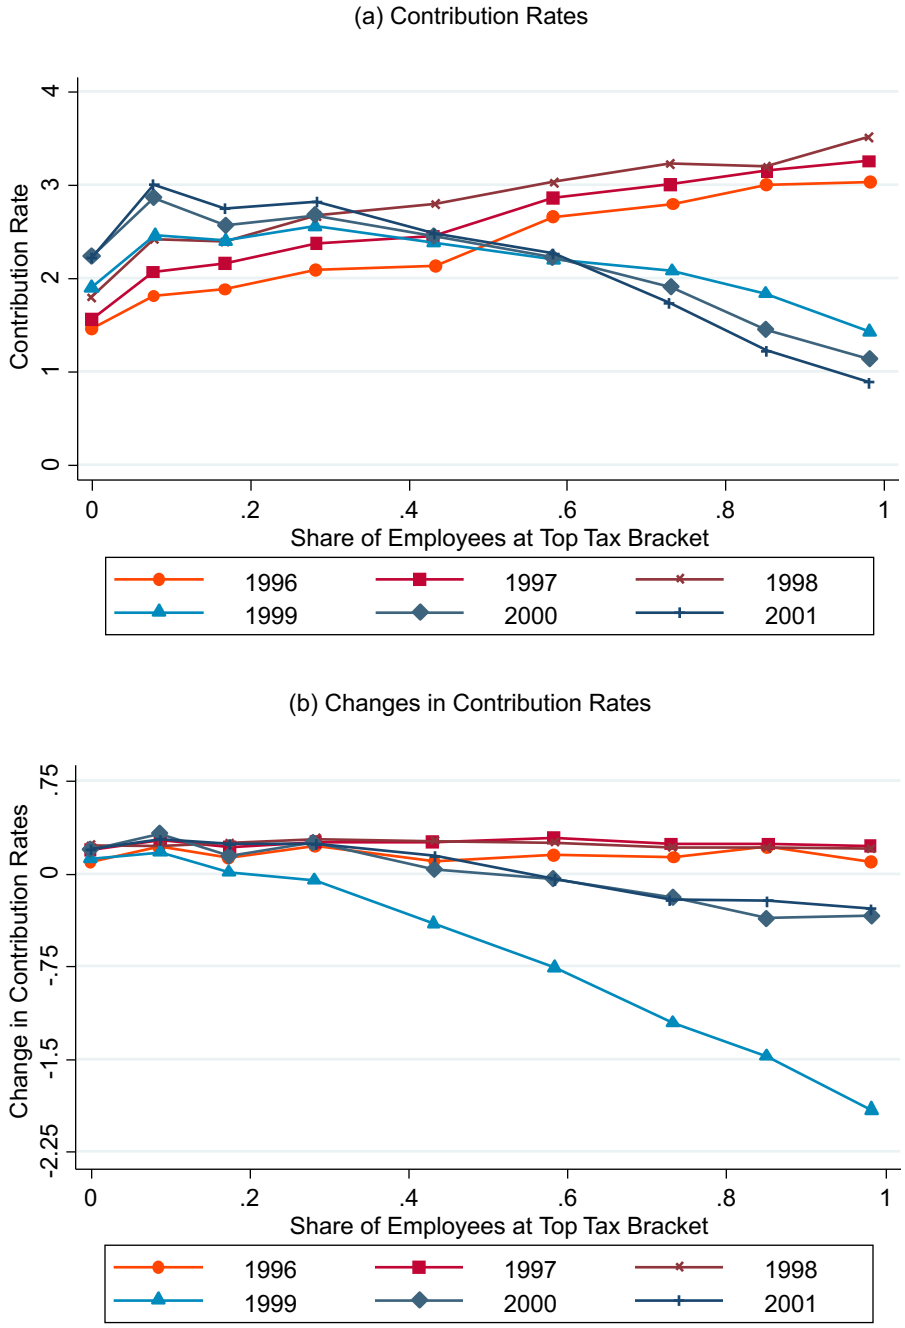

Notes: These figures plot employers' contributions to capital pension accounts as a function of the share of their employees whose earnings were above the top labor-income tax threshold, for years 1996-2001. Panel A plots employer capital contribution rates (as a fraction of labor income), and Panel B plots changes in employer capital contribution rates from the previous year. The observation units are workplaces, defined as all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and then plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

workers had larger capital contribution reductions in response to the reform. This suggests that employers are indeed responsive to changes in their employees' saving incentives, consistent with the hypothesis that employer-provided pension plans reflect the savings preferences of their particular workforce composition.

Regression Analysis. To quantify the firms' responses to the reform, we estimate regressions of the relationship between the change in employer capital contributions and the fraction of workers above the threshold. This also allows us to test the sensitivity of our results to a flexible set of controls. Our baseline estimating equation is of the form:

$$y_{ft} = \beta_0 + \beta_1 above_{ft} + \sum_{s=1996, s\neq 1998}^{2001} [\beta_s (I_{t=s} \times above_{ft}) + \mu_s] + X_{ft} + \varepsilon_{ft}. \quad (1.1)$$

The outcome variable y_{ft} is workplace f's behavior in time t, i.e., annual outcomes grouped at the occupation-firm level. Our first and main outcome variable is the change in employer capital contribution rates from year $t-1$ to year t. The right-hand side variables include the fraction of employees above the threshold in an occupation-firm-year cell ($above_{ft}$), year fixed effects (μ_s), and year dummies interacted with the fraction of employees above the threshold ($I_{t=s} \times above_{ft}$). In this specification, we omit 1998 as the baseline year, so that all the coefficients β_s are estimated relative to 1998. We choose a specification linear in the fraction of affected employees, since Panel B of Figure 1.1 revealed an approximately linear relationship between the change in employer capital contributions and the share of the workforce above the threshold. The main coefficient of interest is β_{1999}. This coefficient captures the relationship between the change in annual employer capital contributions and the fraction of workers above the threshold in 1999 compared to that in 1998, thus estimating the effect of the reform on this relationship. The vector X_{ft} includes various sets of controls, which we add in order to verify the robustness of the estimated effect, as the share of employees above the threshold may be correlated with other characteristics of the firm that may affect the change in contribution rates.

Panel A of Table 1.2 reports the coefficients on the share of employees above the threshold and the interaction of this share with indicators for years 1996 through 2001 (omitting 1998) in regressions that include various sets of controls. In all columns, we include year fixed effects

and cluster standard errors at the workplace level.[19] We multiply the coefficients by 100 to convert them into percentage-point units. Column (1) estimates the baseline regression and in columns (2) to (4) we successively add controls to the vector X_{ft}. Importantly, we add high-order polynomials of the mean workplace-level income, separately for workers below and above the top tax threshold, as well as their interactions with the year dummies.[20] This allows us to further isolate the relationship between employer responses and whether employees are exactly above or below the threshold, by adding an underlying flexible continuous relationship between employer behavior and average labor income.[21] The additional controls that we include are the number of workers in a workplace and its square, as well as their interactions with year indicators, workplace (i.e., 2-digit occupation-firm) fixed effects, and 2-digit occupation-year fixed effects.

[19]Clustering at the firm level instead of the workplace level does not change the statistical significance patterns of our results.

[20]The reported estimates are for polynomials of degree five, but the results are robust to higher- and lower-degree polynomials and are available from the authors on request. The decline in the number of observations from specification (1) to (2) is due to the inclusion of controls for average income separately for employees above and below the top bracket, which excludes workplace-year observations in which all employees are either above or below the threshold.

[21]These controls alleviate concerns, for example, that "good" firms with higher wages may be more likely to respond to the reform and also have a higher fraction of workers above the top tax threshold. We estimated regressions that add controls for percentiles of the workplace's distribution of employee income and found similar results. The analysis is available from the authors on request.

Across all specifications the results are very stable and are in accordance with the graphical results. There is no meaningful relationship between changes in employer capital contributions and the share of employees above the top threshold prior to the reform in years 1996-1998. However, the coefficient on the fraction of employees above the top threshold interacted with 1999 is approximately -2.2 pp and statistically significant at any conventional significance level. Focusing on the specification of column (4) with the full set of controls, this implies that in 1999 employers in workplaces with 100 percent of employees above the top-income tax threshold decreased their capital contribution rate by an average of 2.18 pp more than employers in workplaces with 0 percent of employees above the threshold. For years 2000 and 2001, the coefficients on the fraction of employees above the top threshold are -0.59 and -0.55, respectively, and statistically significant. These patterns are consistent with firms responding substantially just after the reform took place, with a small degree of delayed or gradual responses by some firms.[22]

The stability of the estimated effect across the different regression specifications suggests that the estimated relationship is not driven by omitted variables. Still, a major possible concern is that these results are due to employee responses rather than employer responses. This concern stems from the fact that individuals whose employers contribute to both annuity and capital accounts have the ability to choose a different distribution between capital and annuity contributions than the default set by the employer. In order to address this concern, we conduct a variety of tests detailed in Appendix A.1. First, instead of workplace-level medians, we calculate workplace-level modes, which are not affected by specific individuals, to identify employer contributions rates and find very similar results. Second, we plot the distribution of the distance between employee-level capital contribution rates and the workplace-level aggregates and find that most employee contributions bunch exactly at these aggregates, supporting our method for identifying employer behavior. Third, we find similar results when we focus on the sample of workplaces whose default annuity contribution rate in the years prior to the reform was zero, and whose employees, therefore, did not have discretion in allocating contributions across different types of employer-sponsored accounts. Fourth, since medians more accurately identify default contribution

[22]We find that for workplaces with more than 50 percent of workers above the top threshold, approximately one half of the decrease in employer capital contributions after 1999 is delayed and attributable to firms that did not respond in 1999, and the other half is gradual and attributable to additional responses by firms that responded in 1999.

Table 1.2: *Changes in Employer Contribution Rates to Pension Accounts by the Share of Workers above the Top Tax Threshold*

Dependent Variable:	Δ Capital Contributions			
	(1)	(2)	(3)	(4)
Fraction of Employees Above Top Tax	-0.052*	-0.113*	0.122	0.132
Threshold (Baseline Year 1998)	(0.032)	(0.067)	(0.137)	(0.145)
Fraction of Employees Above Top Tax Threshold Interacted with:				
Year 1996	0.035	0.101	0.202	0.273
	(0.046)	(0.098)	(0.139)	(0.167)
Year 1997	0.061	0.075	0.071	0.096
	(0.046)	(0.096)	(0.126)	(0.147)
Year 1999	-2.126***	-2.167***	-2.181***	-2.182***
	(0.061)	(0.129)	(0.176)	(0.203)
Year 2000	-0.606***	-0.569***	-0.556***	-0.593***
	(0.048)	(0.104)	(0.149)	(0.173)
Year 2001	-0.558***	-0.598***	-0.590***	-0.553***
	(0.046)	(0.101)	(0.152)	(0.176)
Year Fixed Effects	X	X	X	X
Income and Workforce Size Controls		X	X	X
2-Digit Occupation-Firm Fixed Effects			X	X
2-Digit Occupation-Year Fixed Effects				X
Number of Observations	84,764	60,643	60,643	60,643
Number of Clusters	26,775	20,642	20,642	20,642
Panel B				
Dependent Variable:				Δ Annuity Contributions
Fraction of Employees Above Top Tax Threshold Interacted with Year 1999				1.962***
				(0.195)
Panel C				
Dependent Variable:				Δ Overall Employer Contributions
Fraction of Employees Above Top Tax Threshold Interacted with Year 1999				-0.22
				(0.217)

*Notes:*This table reports estimates of employers' responses to the reform as a function of the share of their employees whose earnings were above the top labor-income tax threshold (equation (1.1)). In Panel A, the outcome variable is the change in employer capital contribution rates from the previous year. We regress this outcome on the fraction of workers above the top tax threshold, year fixed effects, the fraction of workers above the top tax threshold interacted with year fixed effects, and different sets of controls as indicated in the table. The coefficient on the fraction of employees above the top tax threshold interacted with other year indicators estimates this relationship relative to the relationship in the baseline year. Income controls include a fifth-order polynomial of the mean workplace-level labor income, separately for workers above and below the top tax threshold, as well as their interactions with year indicators. Workforce size controls include the number of workers in a workplace and its square, as well as their interactions with year indicators. In Panels B and C, we replicate the specification with the full set of controls from column (4) of Panel A, but where the outcome variables are changes in employer annuity contribution rates and changes in overall employer contribution rates to both capital and annuity accounts, respectively. Standard errors are clustered at the workplace level. Coefficients are multiplied by 100 so that they are converted to percentage point units.

rates in large workplaces and since, conceptually, firms are probably more likely to tailor defaults to groups of employees of similar occupations in larger workplaces, we study the robustness of the results to varying the minimal size of workplaces that we include in the analysis. Again, we show that the findings remain qualitatively similar. Overall, the analysis in Appendix A supports our conclusion that the results are driven by firm responses rather than by individual responses.

In addition, in Appendix A.3 we demonstrate that the employer responses in contribution rates were attributable to changes in capital contributions (that is, the numerator) rather than changes in labor income (the denominator). Finally, in Appendix A.4, we show that inadvertently including workers who were covered by collective bargaining likely only attenuates our results.

In sum, our analysis is consistent with the notion that employers design pension plans to reflect the savings preferences of their workforce, and that they respond immediately to changes in their employees' incentives. In the next section, we gauge the magnitude of the employer responses that we estimated.

1.3.2 Employer vs. Individual Responses

In the analysis above, we showed that the average response of employers was large relative to their baseline contribution rates to capital accounts. In this section we assess the magnitudes of the employer responses to the reform by comparing them to the responses of individuals within their private accounts.

The ideal experiment that compares individuals' savings behavior and employers' savings behavior on the individuals' behalf would randomly assign savings decisions to either individuals or their respective employers. To mimic this experiment, we exploit the Danish setting that provides us with administrative records of employee-level savings contributions to both employer-sponsored 401(k)-like accounts and private IRA-like accounts that are managed by the individuals themselves. We focus the analysis only on those who were directly affected by the reform – that is, employees at the top bracket of the labor income tax schedule – and compare their responses in private accounts to those of their employers in their employer-sponsored accounts.

In Figure 1.2, we divide the sample of affected workers into equal-sized groups according to

the share of workers above the top threshold in their workplace. Panel A plots the change in the default contribution rate to employer-sponsored capital accounts (defined only for workers affected by the reform), while Panel B plots the change in the average contribution rate to private capital accounts.

One key difference between employer and individual responses is that for any fraction of employees at the top bracket, the decrease in employer contributions to capital accounts was larger than the individuals' responses. The latter is at most a decrease of 0.75 pp, while the smallest decrease in employer contributions is more than 1 pp. This suggests that most of the overall decrease in capital contributions due to the reform was attributable to employer, rather than individual, responses. It is, in part, due to the fact that baseline contributions to capital accounts in the years prior to the reform were much smaller in individual accounts compared to employer-sponsored accounts (see Table 1.1). Another noticeable difference between the two panels of Figure 1.2 is that there is no gradient in private accounts with respect to the share of employees above the top threshold, while there is a pronounced gradient in employer-sponsored accounts. This suggests that the response of employers in capital pension plans was not crowded out by individual responses in private plans.

To account for differential baseline contribution rates to private vs. employer-sponsored accounts and to understand better what underlies the aggregate responses, we analyze in Appendix A.2 the changes in the two types of accounts at the employee level by studying their respective distributions.

Altogether, the comparison of responses between employer-sponsored accounts and private accounts reveals that, at the aggregate, the effect of the reform on capital contributions was mostly driven by employers. This underlines the large relative role of employers in individuals' overall savings for retirement. Next, we study whether the decrease in employer contributions to capital accounts translated into a decline in overall savings or into increased contributions to other substitutable accounts.

Figure 1.2: *Employer vs. Individual Contributions to Capital Pension Accounts of Workers above the Top Tax Threshold*

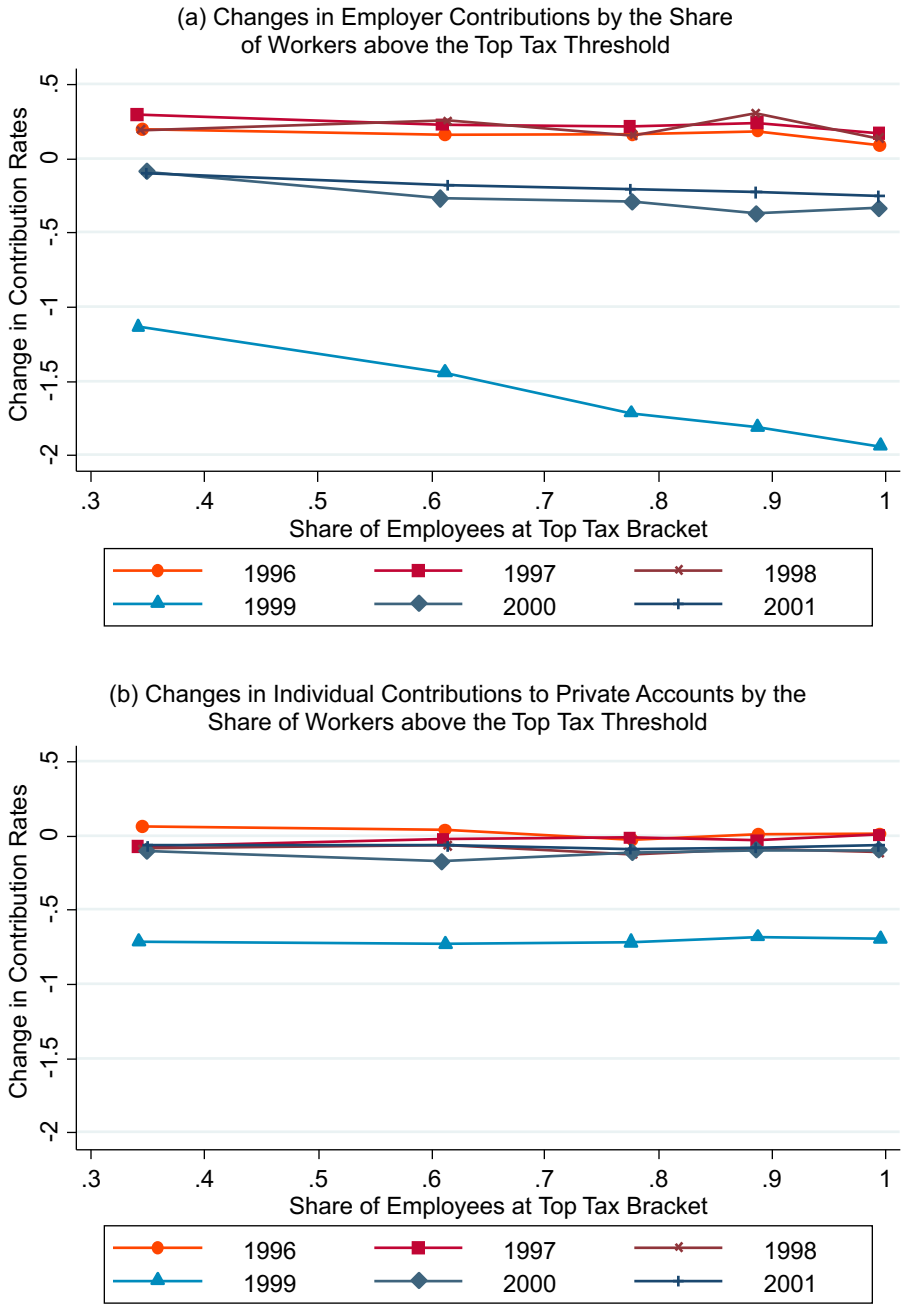

Notes: These figures plot changes in workplace-level contribution rates to capital pension accounts, only for employees with labor income at the top tax bracket, as a function of the workplace's share of employees above the top tax threshold, for years 1996-2001. Panel A plots changes in median capital contribution rates to employer-sponsored (401(k)-like) accounts, and Panel B plots changes in average capital contribution rates to private (IRA-like) accounts. The observation units are workplaces, defined as employees with the same 2-digit occupation code in the same firm. We plot these figures by dividing the sample into equal-sized groups according to the share of employees in the workplace above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals, workplaces with less than five employees, and employees with earnings below the top tax threshold.

1.3.3 Substitution into Annuity Accounts

The empirical analysis of employer responses to the reform has focused so far on capital contributions. To understand the effects of the reform on the overall employer-sponsored savings portfolios of employees, we proceed by analyzing how employer contributions to annuity accounts may have changed.

The reform's decrease in subsidies to contributions to capital accounts had two main effects on employees' savings incentives for workers in the top income tax bracket. As it made contributions to savings accounts less attractive, the reform caused a negative income effect that pushed toward lower levels of total pension savings. At the same time, the reform created a substitution effect due to the decrease in the relative price of contributions to annuity accounts. The relative forces of these two effects determine whether and to what extent employers responded in their contributions to annuity accounts and in their employees' total compensation.

Figure 1.3 plots changes in employers' contributions to annuity accounts. We begin with Panel A, which plots changes in employer contribution rates to annuity vs. capital accounts by year. This graph shows that the time series of changes in employer contributions to annuity accounts essentially mirrored the changes in capital contributions. Moreover, Panel B of Figure 1.3, which replicates Panel B of Figure 1.1 but with changes in employer annuity contributions as the outcome variable, reveals that employer responses in annuity accounts also mirrored the responses in capital accounts as a function of the share of employees above the top threshold. Before the reform, annual changes in both annuity and capital accounts were uniform across workplaces with different fractions of workers above the top tax threshold. In 1999, in response to the reform, employers decreased their capital contributions and increased their annuity contributions as a function of the share of their affected employees and in similar magnitudes. Panel B of Table 1.2 estimates equation (1.1) with a full set of controls and with changes in employer annuity contributions as the outcome variable. We find that in 1999 employers in workplaces with 100 percent of employees above the top-income tax threshold increased their annuity contribution rate by an average of 1.96 pp more than employers in workplaces with 0 percent of employees above the threshold, alongside the decrease of 2.18 pp in capital contributions. In fact, studying the sum of these responses in Panel C of Table 1.2, the evidence is consistent with full compensation of the

decrease in capital contributions by an increase in annuity contributions, so that the change in overall employer contributions is not statistically different from zero.

In sum, the results suggest that the response at the employer-level was driven by a substitution effect, so that the decrease in capital contributions was almost-fully compensated for by an increase in annuity contributions, with no statistically significant effect on overall employer contributions. These employer responses are in contrast to individual responses in private accounts, for which Chetty et al. (2014) find a shift of 57 cents to annuity accounts for each DKr that individuals would have contributed to capital accounts. In Appendix A.3, we additionally show that there was no average effect on labor income, so that we do not find evidence that the potential income effect of the reform was offset by higher wages.

1.3.4 Optimality of Firm Responses

A comprehensive analysis of the optimality of firms' responses would require estimating workers' preferences, modeling their life-time budget constraint, and characterizing their optimal allocations, which is beyond the scope of this paper. In this section we attempt to provide a back-of-the-envelope assessment of how close firms' responses were to "optimal" from the perspective of the employees. Therefore, it is necessary to hypothesize or impute how optimizing individuals would respond to the reform absent any adjustment or information costs. To do so, we choose the actions taken by "attentive" affected individuals, who made changes to their self-managed IRA-like savings accounts in response to the reform, as a baseline benchmark for employee-level optimal response. We choose the response of attentive individuals rather than all individuals since most individuals are passive in their savings behavior, likely due to re-optimization costs (Chetty et al. 2014)[23]

Consider individuals at the top income tax bracket who were directly affected by the reform. In the year of the reform, almost all individuals among this group, who had positive individual capital contributions in the previous year and actively changed their contribution levels, chose to

[23]This is in the spirit of Bernheim and Rangel's (2009) choice-based approach to welfare and its application by Chetty, Looney, and Kroft (2009), who assume that optimal responses to changes in tax rates in the presence of optimization frictions can be recovered by choices when taxes are salient.

Figure 1.3: *Employer Contributions to Annuity Pension Accounts*

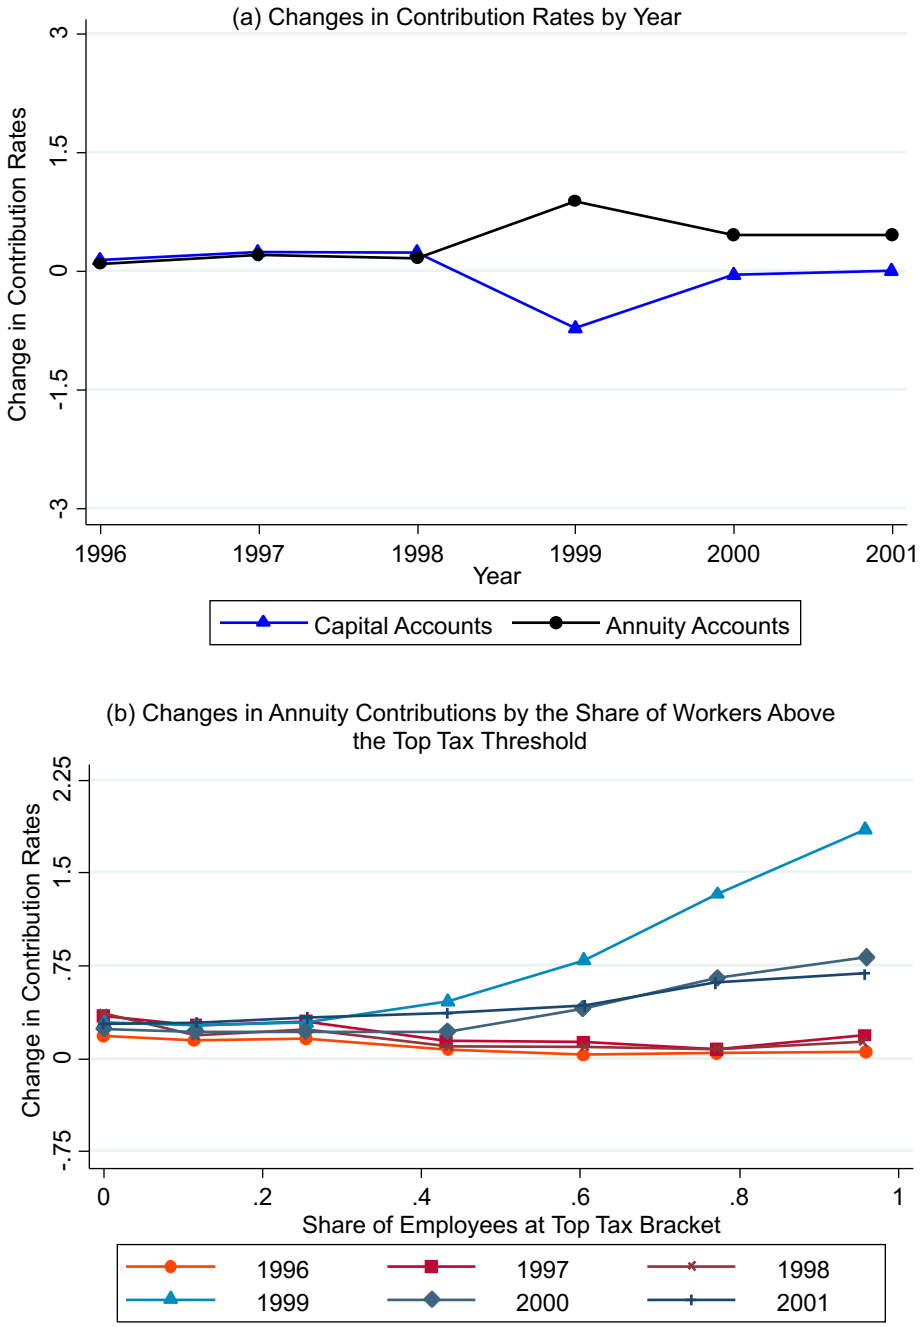

Notes: These figures plot changes in employers' contributions to annuity pension accounts, for years 1996-2001. Panel A plots changes in employer contributions by year, comparing annuity contributions (black line and circles) to capital contributions (blue line and triangles). Panel B plots changes in employers' contribution rates to annuity pension accounts as a function of the share of their employees whose earnings were above the top labor-income tax threshold, for years 1996-2001. The observation units are workplaces, defined as all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

completely exit their individual capital accounts (see Chetty et al. 2014 and Appendix A.2). To the extent that their attentiveness is not systematically correlated with their underlying ranking of saving choices, their actions may be viewed as a crude benchmark for optimal responses of workers above the threshold.[24] We use this individual-level optimality benchmark to characterize whether firms optimally responded to the reform.

We simplify the analysis by focusing on workplaces whose entire workforce is above the top income tax threshold and was directly affected by the reform. Inclusion of workplaces with a heterogeneous mix of workers requires additional assumptions regarding how firms weight the utility of its individual workers, which we choose to abstract from here. In addition, we constrain the analysis to workplaces whose default contribution rate to capital accounts was positive in 1998, the year prior to the reform, in order to focus the analysis on employers that could adjust these contributions downward in response to the reform.

To study the optimality of the firms' responses, Figure 1.4 plots the time series of the fraction of workplaces with positive default capital contributions among those who had positive default contributions in 1998. Within the framework of efficient compensation models, one may expect larger firms to be more responsive to changes in economic incentives if there are significant returns to scale in managing employees' pension products, leading to a more optimal response. In line with this conjecture, we divide workplaces into two size categories: "small" workplaces with 5-19 employees (green line and squares), and "large" workplaces with 20 employees or more (blue line and x's). Following the reform, for both categories of workplace size, we see a sharp drop in 1999 in the fraction of employers who offer capital savings accounts, with continued declines in the following years. In addition, the figure shows that the decreases are significantly larger for larger workplaces, consistent with the notion of returns to scale in managing employee

[24]This is more likely to be a valid benchmark if the financial incentives within the reform led to similarly desirable re-allocations of savings for attentive and inattentive top-income earners. Note, however, that this remains a valid benchmark even if attentive individuals differ by their disutility from making an active choice so long as it is uncorrelated with their optimal allocations. One case where we may worry about this sort of correlation is when the optimal savings adjustment in response to the reform for the "inattentive" affected individuals involves only a small potential utility gain combined with some non-negligible adjustment costs. In this case, they may not respond since for them the costs of re-optimization outweigh its benefits. However, Appendix Figure 1.4 reveals that the difference between the distribution of individual-level annual changes before and after the reform is not distributed smoothly but rather almost entirely concentrated at the mass point of completely opting out. If adjustment costs are distributed smoothly, these findings suggest that opting out of capital accounts is the optimal response for inattentive individuals as well. Nevertheless, the exercise here is merely a rough assessment of optimality and should be taken as such.

Figure 1.4: *Extensive Margin Responses in Capital Pension Accounts by Year*

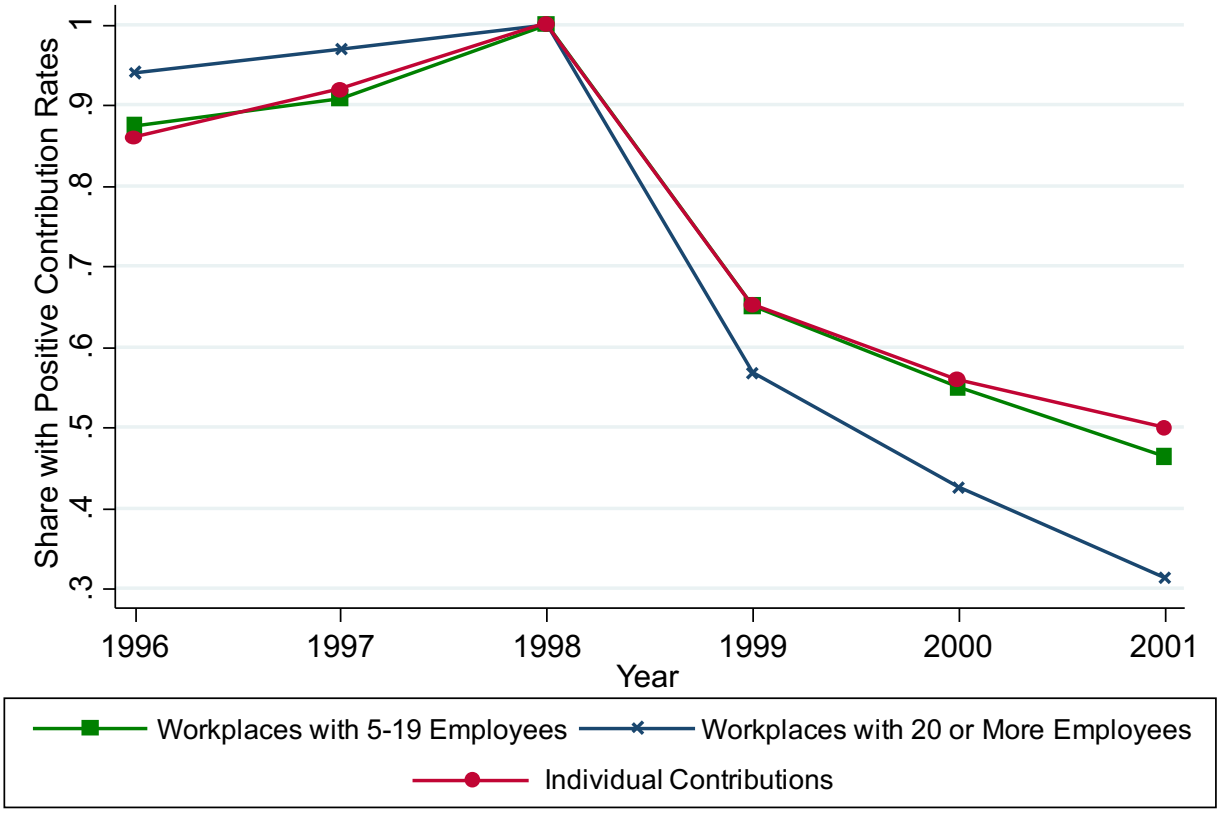

Notes: The green line and squares and the blue line and x's plot the time series of the fraction of workplaces with positive median contributions to employer-sponsored capital accounts among those who had positive median contributions in 1998. We divide workplaces by the size of their workforce, where a workplace is defined as the group of all employees with the same 2-digit occupation code in the same firm. We include only workplaces whose entire workforce was above the top income tax threshold and was therefore directly affected by the reform. In addition, we constrain the sample to private-sector firms and white-collar occupations, and exclude self-employed individuals and workplaces with less than five employees. As a comparison, the red line and circles plot the time series of the fraction of individuals with positive private capital contributions among those who had positive private contributions in 1998. We include only individuals with income levels above the top income tax threshold who were directly affected by the reform. We additionally constrain the sample to employees with white-collar occupations in the private sector, and exclude self-employed individuals.

savings. Specifically, 45 percent of large firms opt out of capital contributions in 1999, compared to 35 percent of small firms. By 2001, almost 70 percent of large firms have opted out of capital contributions, while approximately 52 percent of small firms have opted out by then.

Note that while this response is still far from the assumed optimal benchmark, even for large firms, one important question for policy is how it compares to the average individual response. To answer this question we augment the comparison of employer and individual responses of Section 1.3.2 by focusing here on the "extensive" margin of private capital contributions. This will allow a comparison that is in line with the optimality benchmark of completely opting out of

these accounts. We do so by adding to Figure 1.4 the time series of the fraction of individuals with positive private capital contributions among those who had positive contributions in 1998 (red line and circles). The figure reveals that following the reform approximately 35 percent of individuals who had positive contributions to privately-managed accounts chose to opt out of them, with an overall opt-out rate of approximately 50 percent by 2001. Interestingly, it is apparent in the figure that the response of individuals closely follows that of smaller firms. This is consistent with smaller firms having larger average re-optimization costs per employee that are more similar to the associated costs for individuals than to those for larger firms.

Put together, if opting out of capital accounts is the optimal allocation of savings for affected workers, the exercise above suggests that larger firms make more optimal decisions compared to smaller firms. In addition, the analysis is supportive of the notion that delegating decisions to employers may be welfare increasing, in particular in large workplaces whose size may render managing employee savings less costly.

1.4 Conclusion

This paper provides evidence that employers set contributions to pension savings accounts in accordance with the savings preferences of their workforce, and that they respond immediately and significantly to changes in their employees' savings incentives. In particular, we find that the change in employer capital contributions in response to an increase in their relative price within the 1999 reform was strongly related to the fraction of workers who were above the top tax threshold and were directly affected by the reform. We also find that employers adjusted their employees' overall savings portfolios by significant shifts into the more subsidized annuity accounts, with almost no leakage of overall savings.

Since employer contributions and defaults are extremely effective at increasing individuals' total level of savings, some governments are considering implementing policies that incentivize employer-based savings accounts and default contribution rates. Given the increasing reliance of individual retirement savings on employers' contributions, our findings are promising and encouraging preliminary evidence that they are set in accordance with workers' savings preferences.

However, there are other important aspects of firm responses that we did not address in this paper. For example, our results do not reveal whether firms' behavior is attributable to benevolence or to competition. Additionally, we are unable to evaluate the optimality of responses for firms with heterogeneous workforces, which include a large share of all firms. We believe that addressing these issues is a fruitful direction for future research, as they have potentially important implications for the optimal design of employer-based retirement savings policies.

Chapter 2

The Effects of Physician Prescribing Behaviors on Prescription Drug Use and Labor Supply: Evidence from Movers in Denmark[1]

2.1 Introduction

Much evidence indicates that health is a critical determinant of labor supply (e.g. Currie and Madrian 1999, Cai and Cong 2009, Gaskin and Richard 2012, Cai, Mavromavas, and Oguzoglu 2013, Dobkin et al. 2016, Fadlon and Nielsen 2016). In the United States, over one half of the prime age males not in the labor force have a debilitating condition (Krueger 2016). Gaskin and Richard (2012) estimate that the loss in productivity due just to pain is approximately $300 billion a year. Furthermore, this is increasingly a problem: from 1999 to 2013, mid-life morbidity (including self-reported declines in mental health, and increases in chronic pain and inability to work) among white Non-Hispanics in the United States increased after decades of improvement (Case and Deaton 2015).

Applications to disability insurance suggest that musculoskeletal and mental health disorders

[1]Co-authored with Torben Nielsen

are the most prominent conditions that keep individuals from working (Maestas, Mullen, and Strand 2013). Such conditions are often treated with prescription medications. For example, one half of U.S. men whose health prevents them from working take pain medication on a daily basis (Krueger 2016). There is, however, substantial variation in how physicians treat musculoskeletal and mental health conditions, and it remains unclear how variation in treatment affects labor supply.

In this paper, we use variation in physicians' prescribing rates of medications used to treat debilitating conditions to estimate the impacts of physician prescribing behaviors on patient labor supply.[2] In particular, we study how quasi-experimental changes in the physician prescribing rates experienced by an individual associated with a geographic move affects that patient's prescription drug use and labor market outcomes.

Although the medical literature shows that the use of approved prescription drugs can improve health in the short run, much less is known about whether such prescription drug use can translate into sustained improvements in labor market outcomes. Even the direction of impacts of prescription drug use on labor market outcomes is potentially ambiguous. For example, opioids may effectively treat pain and allow individuals with chronic pain to work, but they also are highly addictive and opioid abuse may negatively influence long-term labor productivity and health.[3] Opioids offer a particularly stark contrast between short run relief and long run consequences, but other prescription drugs may generate similar trade-offs.[4]The effect that physician prescribing rates of drugs have on labor supply remains a difficult question to answer due both to the sparsity of data that connects physicians, prescription drug use, and labor market outcomes, and the potential endogeneity of physician choice.

To estimate the effects of physician prescribing rates, we use Danish administrative data, with which we can link an individual's labor supply, prescription drug use, general practitioner, and residence for almost two decades. Our identification strategy exploits a quasi-random separation

[2]This is a policy relevant question, given that most policies to change individual's prescription drug use would likely work through altering physician's prescribing behaviors.

[3]Opioid abuse is a substantial global problem: 15 million people worldwide suffer from opioid dependence (WHO).

[4]Benzodiazepines are used to treat short-term anxiety and insomnia, however, they can lead to physical dependences and adverse psychological and physical effects, as well as lead to overdoses (Ashton 2005).

of an individual from their general practitioner due to the patient moving to a new municipality. Because the choice of the new physician may be endogenous to the patient's health, we use the pre-move physician's prescribing rate to instrument for the change in physician prescribing rates. Since individuals with prior doctors with different prescribing rates may have different trends in their health (for example due to the effects of their doctor), we match individuals who move to similar individuals with similar pre-period physician prescribing rates who do not move and assign them a placebo moving year. We then take the triple difference in outcomes between the movers and non-movers, the individuals with a high predicted change in physician prescribing rates and with a low predicted change, and after the move minus before, to identify how different changes in prescribing rates affect individual's own prescription drug usage and their labor supply.

The main assumption for this strategy is that other determinants of the relative change in prescription drug use and labor force participation between movers and non-movers are unrelated to the pre-move physician's prescribing rate, outside of its effects on the change in physician prescribing rates. We discuss and address the main threats to this assumption later in the paper and find that they do not bias our results. Our analysis focuses on 730,000 individuals in Denmark aged 30-70 who move across municipalities from 1995-2013, as well as a matched sample of 1,530,000 individuals who do not move. We evaluate the effects of physician prescribing rates for the four drugs that are the most widely used to treat musculoskeletal and mental health disorders: opioids, anti-inflammatories, anti-depressants, and anxiolytics.

Our analysis provides two sets of core results. The first set of results estimates how physicians affect patient's prescription drug use. We find that moving to a doctor that prescribes 1 percentage point (pp) more opioids leads to an increase in opioid usage of .45 pp, with a standard error of .04 pp. Equivalently, moving to a doctor with a one standard deviation higher prescribing rate of opioids leads to a .8 pp (12%) increase in individual's opioid prescription drug use. We find similar sized effects of the prescribing rates of other prescription drugs. This suggests that the choice of physician has substantial effects on the treatment that individuals receive. The second set of results estimates how physician prescription drug rates affect labor market outcomes. We find that for every 1 pp increase in their physician's opioid prescribing rate, an individual's labor

income rank decreases by .12 percentiles.[5] Decomposing this effect, we find there is a significant .15 percentage point decrease on the extensive margin, and evidence which suggests that half of the labor income rank effect is due to effects on the intensive margin. On the other hand, none of the prescribing rates of the other drugs we analyze have consistent discernible effects on the labor outcomes we examine.

The Danish administrative data allows us to look at the labor supply effects along various dimensions. In particular, we identify the effects on individuals' labor income earnings, whether they are on sick leave or receive disability insurance, and their labor force participation. In Denmark, individuals are well insured for health income risks, and are easily able to adjust their labor supply - both downward, by taking sick leave or applying for disability insurance - and also upward, by seeking new employment. This suggests that if physician prescribing rates impact labor market potential, we should be able to see the response in the labor supply measures in the Danish data. Additionally, for most of the drugs we study, Denmark has similar trends in prescription drug use to the United States and other developed countries.[6] In particular, Denmark has a relatively high level of prescription opioid consumption, which has also increased over this time period, much like the United States.

Our methodology builds on other papers that use separations from different types of entities (e.g. places, industries, or firms) to estimate the causal effect of that entity on individual outcomes (e.g. Gibbons and Katz 1992, Chetty and Hendren 2015, and Finkelstein, Gentzkow, and Williams 2016)). Most similar is Finkelstein, Gentzkow, and Williams (2016), who estimate the causal effect geographical locations have on health care spending based on how individual's health care spending changes after they move to a place with a different average spending level.

Our findings contribute to a couple of related literatures. Our paper finds that approximately 30% of the total variation in prescription drug rates is due to causal physician effects. This contributes to a literature that studies the causes of variation in physician practices, and finds that 33-66% of the variation in various physician practices comes from causal physician effects (Davis,

[5]In our analysis, we focus primarily on the effects on labor income rank because it includes both the extensive and intensive margin effect. Our results are robust to other measures of labor income as well.

[6]The only exception is of anti-anxieties: in Denmark, the use of anti-anxieties have decreased from 1995-2013, while their use has increased in the United States.

Gribben, and Lay-Yee 2000, Grytten and Sorensen 2003, O'Neill Kuder 2005, Wang Pauly 2005, Mercuri et al. 2012, Cutler et al. 2015, Li Laxminarayan 2015, Currie, MacLeod, and Van Parys 2016, and Molitor 2016).

The estimation of the effects of physician prescribing rates on labor supply adds to a growing literature that estimates the effect of prescription drugs on individual outcomes. An extensive medical literature reviews the effects of medical drugs on health metrics, but often only over short periods of time.[7] Additionally, recent work estimates the effects of different prescription drugs on labor supply. Kilby (2015) finds that a decrease in use of opioids due to the onset of the Prescription Monitoring Program laws increases the number of absent days at work for individuals with workers' compensation injuries and those on short term disability with pain related diagnosis codes.[8] Both B́utikofer and Skira (2013) and Garthwaite (2012) look at the effect of Cox-2 inhibitors (part of the anti-inflammatory class) on employment by looking at the abrupt removal of Vioxx from the market in 2004 and find a small decrease in labor force participation due to the removal of Vioxx.[9]

The paper proceeds as follows: Section 2 details the data we use and the institutional framework in Denmark. Section 3 outlines our empirical strategy. Section 4 estimates the effect of a change in physician's prescribing rate on individual's drug use while Section 5 looks at the effect on labor market outcomes. Section 6 concludes.

[7]In medical control trials, opioids are found to decrease pain in the short term, but there is also a large amount of drop out due to adverse events or lack of efficacy (Shaheed et al. 2016). In terms of long-term opioid therapy, in a review Chou et al. (2015) find that "reliable conclusions about the effectiveness of long-term opioid therapy for chronic pain are not possible due to the paucity of research to date." Anti-inflammatories do decrease short term pain better than the placebo, but have little benefit after two weeks of use, and serious adverse effects associated with them (Van Tudler et al. 2000, Bjordal et al. 2004, Lin et al. 2004). In a review, Dell'osso and Lader (2013) find that the risk benefit ratio for benzodiazepines is positive in short-term use, but it remains unclear whether short-term benefits outweigh the possible risk of dependence. A large amount of evidence finds that anti-depressants are better than placebo on quality of life and pyschosocial outcomes, (Stewart et al. 1988, Ceulemans et al. 1985, Stewart et al. 1993, Kocsis et al. 1997, Lydiard et al. 1997, and Bollini 1999). The one medical control trial that we know of that looks at the effects of prescription drugs on labor supply is a study by Agosti, Stewart, and Quitkin (1991) which randomly assigned an anti-depressant or a placebo and looked at labor outcomes (N=43). It found a negative effect on hours worked in the 6 weeks of follow-up, but it was not statistically significant.

[8]In the Section 6, we discuss the reasons why our two papers may have found different results.

[9]There are a couple of additional papers that look at the effects of prescription drug types that we do not analyze in this paper. Daysal and Orsini (2014) analyze the effect of Hormonal Replacement Therapy on employment of middle-age women using the timing of the release of information of the potential hazardous effects of HRT in 2002 and deduce that HRT increases employment. Thirumurthy, Zivin, and Goldstein (2008) look at the effects of anti-virals on labor supply in Africa. Finally, Currie, Stabile, and Jones (2014), look at the effects of ADHD use on education outcomes in Canada.

2.2 Data and Institutional Framework

2.2.1 Data Sources

The ideal data for an analysis of the impact of physician's prescribing rate on labor supply contains information on individual's physicians, their prescription drug use, and their labor supply. We create this dataset using several Danish administrative data sources to build a panel of individuals from cohorts 1925-1980 covering the years 1995-2015, with information on their prescription drug use, their general practitioner, multiple measures of labor supply, and their geographical history, which is necessary for our particular identification strategy. In our analysis, we restrict our sample to individuals aged 30-70 to cover the working age population.

Prescription Drugs: The prescription drug use data covers every prescription purchased by an individual from the 1920-1980 cohorts from 1995-2015.[10] In our analysis, we focus on the outcome of whether an individual purchases a particular prescription drug in a given year, however, for robustness we also look at whether an individual has a pick-up of a prescription drug within a month.[11] We identify four types of prescriptions using the Anatomical Therapeutic Chemical (ATC) classification system. Opioids are defined as ATC code N02A, anti-inflammatories as M01A, anti-anxieties as N05CD and N05BA (Benzodiazepines and derivatives), and anti-depressants as N06A.[12]

Physicians: The patient database covers every visit and service charged by a physician from 1987-2013. Physicians are categorized by a 2 digit speciality. We restrict our focus to visits to general practitioners (GP) from 1995-2013. For more details on the identification of physicians in the data see Appendix B.1.

Labor Outcomes: We use a variety of labor market information to fully characterize the effect prescription drugs have on economic outcomes. The measures we use are labor income rank, log

[10] In Denmark, over the counter medication maybe prescribed (and thus seen in the data), since this allows some medicines to be subsidized. On the other hand, over-the-counter drugs that are without a prescription and drugs administered at hospitals are not covered in the data.

[11] We do not know whether individuals actually take the prescribing drug; however, we will often refer to a prescription drug purchase as "taking a the drug," with the implication being they purchased the drug with the intention of taking the drug.

[12] We group together all benzodiazepines and derivatives, even though some of them are classified as hypnotics and sleeping aids (N05CD).

labor income, labor force participation, disability insurance receipt, and two measures of sick pay receipt.

Labor Income: We measure labor income as taxable labor and self-employed earnings. For the main outcome variable in our analysis, we convert labor income into percentile ranks within an individual's year of birth, the year, and their gender using the full sample of the Danish population. In the case of ties, we define the rank as the mean rank for the individuals in that group. For example, if 20% of women in the year 2000 who were born in 1950 have zero income, then they would all receive a rank of .10. We follow Chetty et al. (2014) who also use this measure because it creates a measure of relative income that is comparable across ages, years, and gender, that is not overly influenced by the tails but still keeps their cardinal ordering; and that includes individuals with zero labor income.

For additional information on the other measures of labor market outcomes we use see Appendix B.1.

Geographical Information: To identify moves between municipalities, we use information on individual's annual municipality residence. Appendix Figure B.1 plots a map of Denmark, showing the municipality boarders. We classify a move as when an individual lives in two different municipalities in two consecutive years. For each residence, we know the date of first residence, so we are able to identify the exact day of the move.[13]

2.2.2 Institutional Framework

In this section we discuss the institutional setting in Denmark with respect to the health care system and the labor market.

Health Related Institutions: In Denmark, general practitioners act as gate keepers to the primary care system and provide referrals to other specialty care. There is no direct to consumer advertising for prescription drugs, so individuals get the majority of information from their doctors. Both general practitioners and psychiatrists can prescribe psychiatric drugs.

[13]In 2007, Denmark changed municipality boundaries by merging some municipalities. After the reform, the number of municipalities went from 293 to 99. We use the new municipality definitions and construct residence based on the new municipality borders prior to 2007. Due to anonymity concerns, we merge municipalities with fewer than 5000 residents to a larger municipality. All of these municipalities happen to be islands, so we merge them with the municipality that they have ferry access to.

Cost of Prescription Drugs: Unlike most healthcare in Denmark, prescription drugs have some copay. Appendix B.1 details the exact rules, and Appendix Figure B.2 Panel A shows the average price paid by the consumer (in 2015 dollars) for one pick up of each type of drug, while Panel B shows the average total non-subsidized cost for one pick up of each type of drug. It shows that copays for the drugs we study are generally between $5-$25, while the total costs ranges from $10-$85.

Assignment to General Practitioners: General Practitioners in Denmark can have up to 2,500 patients per year.[14] After 1,600 patients, GPs can deny new patients. Individuals in dense areas must choose a GP within 5 km of their residence, and in rural areas they must choose a GP within 15 km of their residence. If they don't actively choose then the municipality assigns them a physician.

Labor Market Related Institutions:

A sick individual in Denmark who is under 65 can decrease their labor supply, and still receive income by taking a temporary sick-leave or applying for Social Disability insurance. All employers are obliged to provide sick-pay benefits, which fully replaces income for at least the first couple of weeks. When the employer's sick-pay benefits end, the local government must provide sick-pay benefits equivalent to the prevailing unemployment benefit rate up until a year after the worker has stopped working. If the worker remains sick and unable to work, he or she can apply at the municipality level for Social Disability Insurance (Social DI) benefits that provide income permanently. The program is moderately generous - for example, in 2000, subject to income-testing against overall household income, a successful application amounted to roughly DKK 127,500 ($16,000) per year.[15]

After turning 60, and before they reach the old-age pension retirement age, individuals who have been members of a voluntary unemployment fund for a sufficiently long period are eligible for the Voluntary Early Retirement Pension (VERP). Approximately 80% of the population is eligible for VERP, which provides an annual income that is 90% of previous earnings, but maxes at

[14]This is relaxed in rural areas where there are fewer physicians per capita, due to these regions having problems attracting GPs to these areas.

[15]While Social DI is state-wide scheme, it is locally administered by regional councils and municipality case workers, which has led to differential rejection rates across municipalities ranging from 7-30% (Bengtsson 2002).

the unemployment benefit level which was 148200 Kroner in 2000 ($18,525). At the full-retirement age of 67 (or 65 for those born after July 1, 1939), all residents become eligible for the Old-Age Pension (OAP), which provides income-tested annuities of up to roughly DKK 87,000 ($10,900) per year (at 2000 rates).[16]

2.2.3 Summary Statistics

Prescription Drug Summary Statistics

In Appendix Table B.1, we present summary statistics for the sample we using in our analysis.[17] For prescription drug use, approximately 6.5% of the sample have an opioid prescription purchase within a year. Alternatively, 19% purchase prescription anti-inflammatories, and 6-7% purchase anti-anxieties, and 7-9% purchase anti-depressants.

Over the time period we study, prescription opioid use and anti-depressant use increased, while anti-anxiety use decreased. Figure 2.1 shows the fraction of individuals who use the prescription drugs we study by year from 1995-2015. Panel A shows that opioid prescription drug use has increased steadily from 1995 (5%) to 2015 (9%). Panel B shows that anti-inflammatory use prevalence increased from 16% use in 1995 to 22% use in 2005, at which point usage peaked, and started to decrease slightly. Anti-anxiety use in Panel C has had a remarkable decrease in usage from 11% to 4% over the time period, while anti-depressant use has increased from 5% to 10% in 2010, at which point it flattened off.[18]

For all the prescription drugs we study, use increases substantially by age and is higher for women than it is for men. This is depicted in Figure 2.2, which shows the fraction of individuals who use the different types of prescription drugs at least once in a year by their age and gender. The largest differences by gender, as a percent of use, are for the mental health drugs. At age 50, women are about 60% more likely to use both anti-anxieties and anti-depressants than men.

[16]Note that for individual's over 60 who have been employed, VERP is more generous than DI. However, if individuals have been long-term unemployed, then they may not be eligible for VERP, or it may not be as generous.

[17]Specifically the mover and the non-mover sample for individuals aged 30-70, from the 1925-1980 cohorts, for the years 1995-2015, for up to three years prior to the move, and for up to three years after the move.

[18]The large decrease in anti-anxieties is likely due to a change in recommendations for prescribing them for shorter periods of time. The formal recommendation change occurred in September 2004.

Figure 2.1: *Annual Any Use of Prescription Drugs by Year*

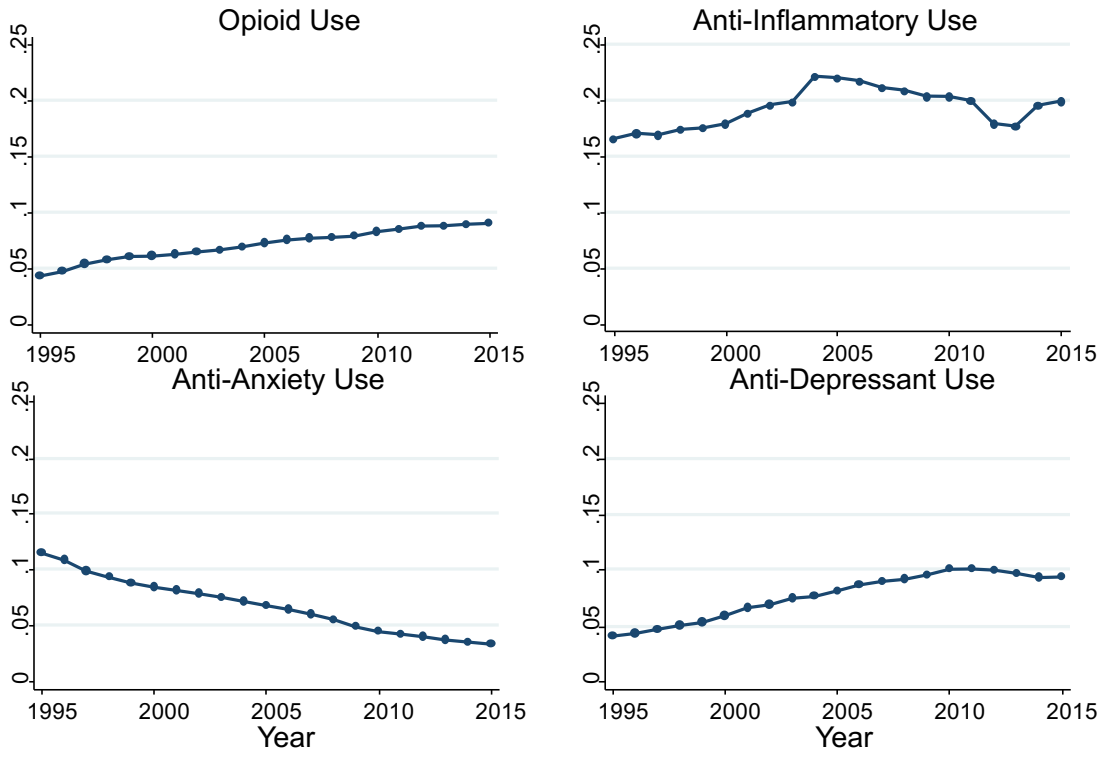

Notes: This figure plots the fraction of individuals who use the four prescription drug classes we study by year. Use is defined by having at least one purchase of that prescription drug up within the year. The full population of individuals aged 30-70 from the 1925-1980 cohorts are included in the estimates. Prescription drug use is classified by the Anatomical Therapeutic Chemical (ATC) classification system. Under this system, Opioids are ATC code N02A, anti-inflammatories are M01A, anti-anxieties are N05BA and N05CD, and anti-depressants are N06A.

Comparison To United States: To understand how Denmark's prescription drug use compares to the United States, we compare available measure of use over the period we study.

Opioids: Prescription opioid use in Denmark is comparable to use in the United States. Using data from International Narcotics Control Board and World Health Organization Population data, Appendix Figure 2.3 plots the annual consumption in Morphine Equivalence Mg/Capital (a measure of opioid use) from 1995-2014 for the United States and Denmark.[19] It shows that prior to 2001, Denmark had higher rates of opioid use, but the United States had much higher growth, so that by 2008, the United States had approximately 40% more consumption in ME mg/capita

[19]This data was aggregated together by the Pain and Policy Studies Group (PPSG) at the University of Wisconsin-Madison. PPSG developed a Morphine Equivalence (ME) metric using conversion factors from the WHO Collaborating Center for Drugs Statistics Methodology for the 6 principal opioids used to treat moderate to severe pain:Fentanyl, Hydromorphone, Methadone, Morphine, Oxycodone, and Pethidine. The ME allows for comparisons between countries of the aggregate consumption of these principal opioids (total ME).

Figure 2.2: *Annual Any Use of Prescription Drugs by Age and Gender*

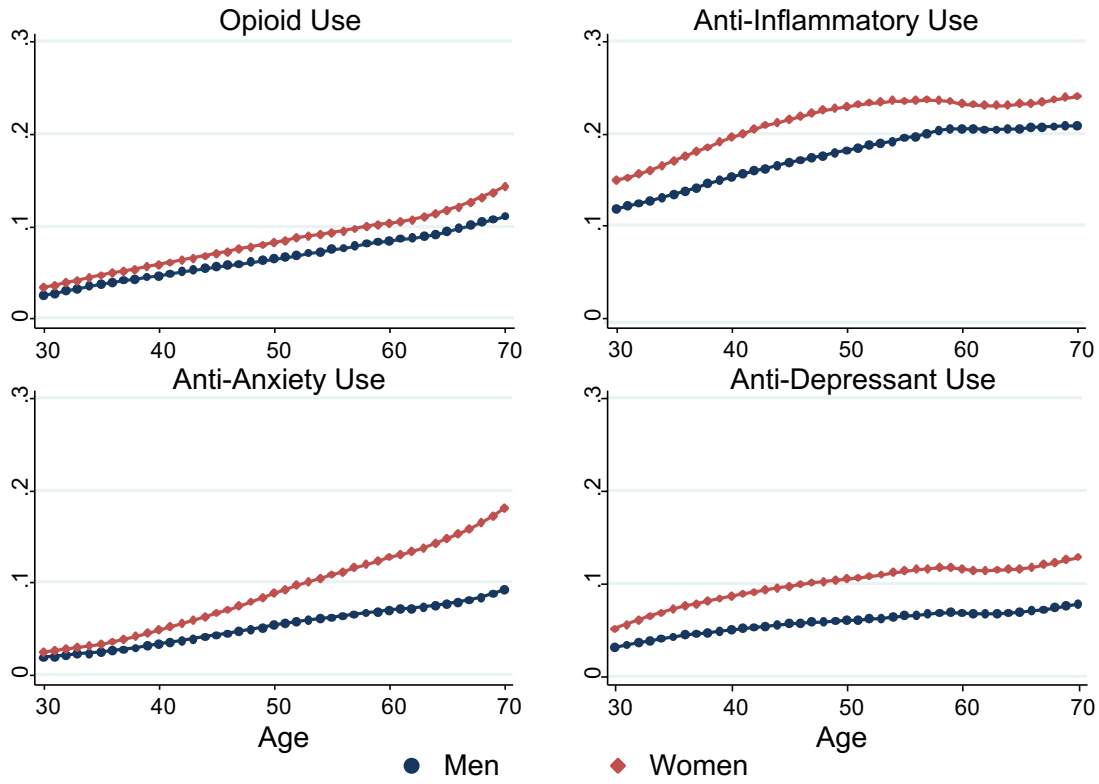

Notes: This figure plots the fraction of individuals use the four prescription drug classes we study by the age and gender of the individual. The average use of women is plotted with red dots and lines, while the average use of men is plotted using blue dots and lines. Use is defined by having at least one pick up within the year. The full population of individuals aged 30-70 from the 1925-1980 cohorts for the years 1995-2015 are included in the estimates. See Figure 2.1 notes for how each type of drug is classified.

than Denmark. Though, compared to other countries, Denmark still has a relatively high usage of opioids.[20]

There is relatively little data in the United States on usage for the other prescription drugs we study. For a comparison, we use biennial information from Kantor et al. (2015)'s analysis of the National Health and Nutrition Examination Survey (NHANES) of 37,959 non-institutionalized US adults aged 20 years and older from 1999-2012. The metric reported is self-reported use in the past 30 days. We compare this to indicators for any annual or any monthly pick ups from 1999-2012 for individuals aged 32-70 using the administrative Denmark data. These are not directly comparable

[20]Only Canada and the United States had higher per capita use on average from 2010-2014.

measures, so we focus on comparing trends rather than levels.[21]

Anti-Inflammatories: The trends for anti-inflammatory use are broadly similar between the two countries. Appendix Figure B.4 Panel A shows that while in Denmark there was a slow increase in anti-inflammatory use up until 2006, and then a slow decrease - in the United States, there was an increase up until 2004 and then a marked decrease (consistent with the removal of Vioxx from the market), after which it stayed fairly steady.

Anti-Anxieties: Appendix Figure B.4 Panel B shows that while in Denmark there was a steady decrease in anti-anxiety medication from 2000-2012, the United States had a slight *increase* in usage. This marks the largest difference in the prescription drug trends for the two countries.

Anti-Depressants: Appendix Figure B.4 Panel C shows that in both Denmark and the United States there has been an increase in anti-depressant use from 2000-2012.

Labor Market Outcomes Summary Statistics

Appendix Table B.1 also shows the summary statistics for the labor force outcomes we analyze in Section 5. Approximately 81-83% of the sample has a positive labor force participation. 8-9% take employer sick pay and 8-10% take municipality sick pay, and 7-8% of the population is on disability receipt.[22]

2.3 Empirical Strategy

In this section, we describe our identification strategy for estimating the effect of physician prescribing rates on individual's prescription drug usage and their labor supply.

We first describe the ideal experiment, then explain our identification strategy and the necessary

[21] Our measures and the measures from the Kantor et al. (2015) analysis differ in meaningful ways for the following reasons: first, our monthly measure is based on pharmacy purchases, whereas their measure is based on self-reported drug use in the past 30 days. If a purchase lasts longer than a month, than this would underestimated the amount of use within the past 30 days. Additionally, since we don't measure individual's actual usage, purchases from the pharmacy might over-estimate actual use. Finally, the distribution of ages of our samples is not the same, though both samples are balanced by age over the time period studied.

[22] We might expect that fewer individuals are on municipality sick pay than employer sick pay, since it kicks in after employers coverage for sick pay ends. However, municipality sick pay also includes pay for paternity leave. While we have attempted to decrease the influence of this by setting it equal to missing by for women who have had a baby within the year or the last year, due to Denmark's generous paternity leave system, this may not cover all individuals who are on sick pay due to paternity leave.

assumptions required for an unbiased estimation. After we explain the general identification strategy, we specify how we implement it.

2.3.1 Identification Strategy

The ideal experiment for answering how physician prescribing rates affect individual's prescription drug use and labor supply would be to randomly assign individuals to physicians with different prescribing behaviors and compare differences in the individual's prescription drug use and their labor supply. Absent this, we propose the following methodology that uses quasi-random changes in physician prescribing rates.

The two main endogeneity problems without random assignment are that individuals may choose their general practitioner based both on current health shocks and their expected future health.[23] To overcome these endogeneity problems, we propose analyzing changes in doctors that are unlikely to be caused by health shocks - changes due to a move between municipalities. This identification strategy relies on the assumption that cross-municipality moves cause a quasi-exogenous separation from an individual's physician that is unrelated to the relationship between their health and their physician's prescribing rate, which we extensively validate below.

While we assume the separation is exogenous, individuals likely choose a new physician based on recent health shocks, which would make their new choice of doctor endogenous. We therefore only use the variation in the individual's prior physician's prescribing rate, which is predictive of the change in physician prescribing rates due to mean reversion in the choice of physicians and their prescribing rates. For example, individuals with a prior physician who has a high prescribing rate will on average have a lower prescribing rate after they move causing a decrease in their physician prescribing rate.

Since individuals with different prior physician prescribing rates may have different health trends, we match each individual to a similar "control" individual who doesn't move at time T based on their age, sex, education, prior physician prescribing rates, and a placebo moving year. We compare their outcomes around the time of the move with the assumption that absent the move, the movers' outcomes would evolve similarly to the non-moving sample with respect to

[23] Individuals would take into account their future health if there is some cost to switching physicians.

the pre-period physician's prescribing rate.[24] Due to a baseline counterfactual separation process of the control non-mover sample from their physician, we use the relative predicted change in physician prescription drug rates between the treatment and the control sample as our measure of the intensity of treatment. Finally, we control for origin by destination by year relative to move fixed effects, to control for location effects which may correlate with doctor prescribing rates.

Our strategy relies on the following assumptions. First, that the pre-period physician's prescribing rate is a strong predictor of the change in physician prescribing rates for individuals after a move. In Section 3.3.2 we show evidence that after a move, physician prescribing rates converge by 75% to the mean, so that the prior physician's prescribing rates are indeed a strong predictive of the change.

The second assumption for our identification strategy is that other determinants of the relative change in prescription drug use and labor force participation between movers and non-movers are unrelated to the pre-period physician's prescribing rate, outside of its effects on the change in physician prescribing rates. There are three main concerns for why this may be violated. First, we are concerned movers and non-movers may have different trends in the relationship between their drug use or their labor supply and their pre-period physician's prescribing rate,. To address this concern, we look at three years prior to the move to show how the relationship between the outcome variables and the instrumented change in the doctor's prescribing rate evolves and find little evidence of differential trends between the movers and non-movers. We also look at the effect of the instrumented change in the physician prescription drug rate on drug use around the time of the move at the *monthly* level and still find no differential pre-trends suggesting that the effects we see are not due to differential simultaneous shocks.

Second, we may be concerned that individuals with different doctors have differential effects from the move aside from the effect of the change in doctor prescribing rates, since individuals with different doctors might move for different reasons. In some specifications, we control flexibly for different effects of the move by age, previous income, and gender and find similar results, suggesting this is not a problem. Additionally, to relax the assumption that other determinants of the relative change in prescription drug use and labor force participation between movers and

[24]This strategy is similar to other papers that use matching strategies - for example J$'$ager (2016) - who matches workers at firms where an employee dies to workers at other firms where an employee doesn't die.

non-movers are unrelated to the pre-period physician's prescribing rate, we use heterogeneity by the distance of the move, which results in variation the intensity of treatment. Conditional on the origin physician's prescription rate, longer distance moves are associated with larger changes in physician prescription drug rates since individuals are less likely to continue to see their old physician after the move. Using this heterogeneity in the treatment allows us to assume instead that other determinants of the relative change in outcomes between movers who *move long distances and short distances* are unrelated to the pre-period physician's prescribing rate. We find similar results using heterogeneity by the distance of move and as main method of comparing movers and non-movers.

Third, we may be concerned that there are other characteristics of the physicians that are correlated with the prescribing rate that are causing the changes in prescription drug use or the changes in labor supply, instead of the physician's prescribing rate. While we cannot entirely rule this out, we control for other characteristics of the physician that are observable - in particular their prescribing rates for other drugs.

Appendix B.2 presents a theoretical model that motivates the empirical identification strategy we use, and Appendix B.3 gives a specific example of our identification strategy by considering two individuals who move from Aarhus to Copenhagen.

2.3.2 Implementation of Identification Strategy

In this section, we describe the essential details for the implementation of our identification strategy: how we assign individuals to doctors, how we identify physician prescription rates, and how we choose the mover and non-mover sample.

Linking Patients to Primary Care Physicians

For each patient-year (it), we link patients to a primary care physician based on the General Practitioner (GP) they saw most in the surrounding 3 years $(t-1, t, t+1)$.[25] This is used for calculating physician prescribing rates. Additionally, we assign movers and non-movers one "pre-period" physician based on the physician they see the most in the three years prior to the

[25]Before we match individuals, we first drop all general practitioners who see fewer than 2000 patients ever, and 400 patients within a year to ensure the GP is in practice throughout the year and sufficiently involved in the health market.

year of the "move," and one "post-period" physician based on the physician they see in the three years after the "move". These physician assignments are used for identifying the predicted change in physician prescribing rates. For more details on how we assign physicians to individuals see Appendix B.4.

Measuring Physician Prescribing Rates

We measure physician's prescribing rates based on the prescription drug use of their patient.[26] We take out variation due to observable and immutable patient characteristics to minimize the amount of variation that is purely due to selection. To do this, we estimate physician fixed effects while controlling for the individual's age, gender, education, and the year. Specifically, we identify:

$$Drug_{it} = p_{j(it)} + \alpha_{af} + \zeta_y + \tau_e + f(age, year, female, highschool) + \epsilon_{it} \tag{2.1}$$

Where $p_{j(it)}$ is a set of fixed effects for each physician j, that is equal to one if individual i has been assigned to physician j at time t, and α_{af} are age by gender fixed effects, ζ_y are year fixed effects, τ_e are fixed effects for 10 different degrees of education, and f is a flexible function of age, the year, an indicator for female, and an indicator for only high school education.

Variance of Physician Prescribing Rates: We estimate that there is substantial variation in the estimated physician prescribing rates. Table 2.1 row 2 reports that the standard deviation for opioid prescribing rates (column 1) is .018 or 1.8pp. For anti-inflammatories, it is 2.9 pp, for anti-anxieties it is 2.4 pp and for anti-depressants it is 1.7 percentage points (anti-depressants).[27] This is a substantial amount of variation: as a proportion to mean usage, a standard deviation in physician opioid prescription rates is: 25%.

Note that for these estimated physician prescribing rates, which do not contain variation from age, gender, and education, still contain variation due to differences in demand based on unobservable health characteristics (selection). When we look at the effect of a change in physician prescribing rates on individual's drug use, we will estimate the fraction of the remaining variation

[26]Note that this includes not only the direct prescriptions made by the physician, but also the prescriptions made by other doctors that the GP referred the patient to.

[27]While our estimates are measured with error, we measure that the signal to noise ratio is about 95%. Therefore the standard deviation of the signal is very close to the estimates.

Table 2.1: *Standard Deviation of Physician Prescription Rates*

	Opioids (1)	Anti-Inflammatories (2)	Anti-Anxieties (3)	Anti-Depressants (4)
Standard Deviation of Phy. Effects				
Raw Physician Averages	0.023	0.036	0.029	0.021
Estimated Physician Effects	0.018	0.029	0.024	0.017
Intstrumented Change in Physician Prescription Rates	0.009	0.015	0.012	0.010
Mean Prescription Rate	0.070	0.193	0.075	0.078

Notes: This table presents the standard deviation of physician prescribing rates, for each type of drug. Prescribing rates are the fraction of patients who pick up a prescription for the drug within a year and can range from 0 (none of the physician's patients take the drug) to 1 (all of the physician's patients take the drug). Individuals are assigned to primary care physicians in a given year based on the general practitioner they see the most in the year before, the year of, and the year after. All individuals aged 30-70 from the cohorts 1925-1980 and from the years 1995-2015 are used to calculate the prescribing rates. Physicians with fewer than 1000 patients are excluded. Row 1 reports the standard deviation of raw physician averages. Row 2 reports the standard deviations of the physician effects we estimate from Equation 1, which takes out variation from the patients age, education, gender, and the year. Row 3 reports the standard deviation of the actual treatment we use, which is the instrumented relative change in physician prescribing rates based on the individual's pre-move physician. The last row reports the fraction of annual usage for individuals aged 30-70, cohorts 1925-1980, and the years 1995-2015. Column (1) reports the standard deviation for opioid prescribing rates, column (2) reports the standard deviations for anti-inflammatory (NSAIDs) prescribing rates, column (3) reports the standard deviations for anti-anxiety (benzodiazapine) prescribing rates, and column (4) reports the standard deviations for anti-depressant prescribing rates.

that is due to demand.

Mover and Non-Mover Sample

We identify the mover (treatment) sample and the non-mover (control) sample based on the individual-year, for the years 1995-2013 and individuals aged 30-70. The sample of mover-years is identified by the year that an individual moves to a new municipality. We create a sample of non-mover-years as a control group so that it exactly matches the treatment sample by the year, their age, sex, education, quartiles of their pre-period physician's prescribing rates for the four drugs we study, and the quartile of their rank of average labor income from $T-8$ to $T-4$. Within the set of all non-mover-years that match the treatment group of mover-years on these characteristics, we take a random sample so that the control group is twice the size as the treatment group, and drop any mover-years that have no control group.[28] The non-mover-year is thus a placebo "move" year for each matched non-mover.

Despite the matching, there are still some differences between the two samples, shown in Appendix Table B.1. The mover sample has slightly higher prescription drug use than the non-mover sample. Additionally, the labor income rank of movers is slightly lower than the non-movers as is other measures of current labor outcomes. These differences in levels are small and do not by itself pose a threat to the design, which compares the trends in the relationship between drug use or labor market outcomes and the relative predicted change in physician prescribing rates. In our analysis in Sections 4 and 5, we find no differences in the pre-trends, which validates our design.

2.3.3 First Stage Effects of Treatment

In this section, we first show that individuals who move are significantly less likely to see their pre-period physician in the post-period than similar non-movers. Second, we show that movers have a larger average change in physician prescribing rates after the "move" conditional on their pre-period physician's prescribing rate. This establishes the first stage of our identification strategy and forms the treatment which we use to identify the effect of prescription drug use on an individual drug use in Section 4 and individual labor supply in Section 5.

[28] In the rest of the analysis, we reweight and control group cells that fewer than twice the size then the treatment group.

Figure 2.3: *Fraction of Individuals that See Pre-Period General Practitioner by Year Since Move*

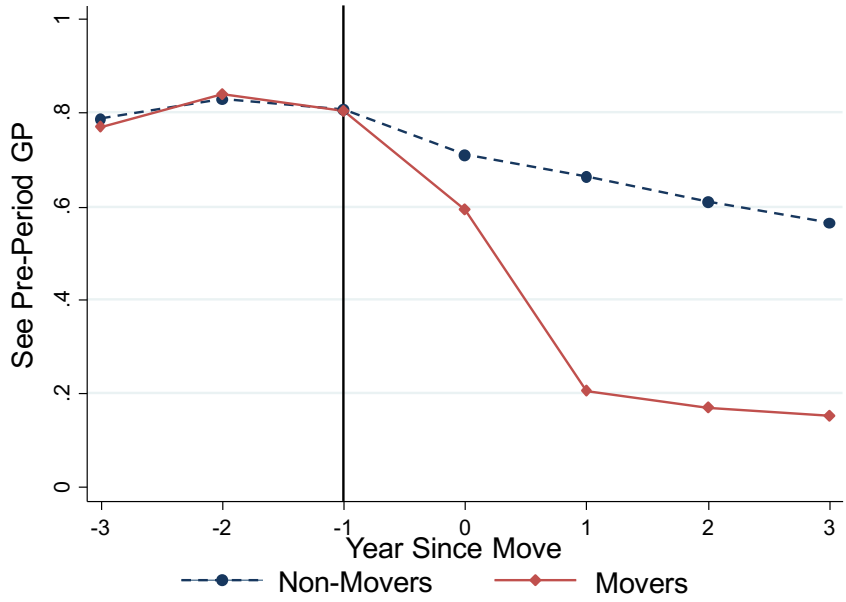

Notes: Figure 2.3 plots probability of individuals seeing their pre-period General Practitioner in each year three years prior and three years after a move between municipalities for movers (in red dots and lines) and a matched control sample of non-movers who are assigned a placebo moving year. The control sample is matched exactly on age, gender, education degree, the year, quartiles of each of the four physician prescribing rates, and a quartile of individual's average labor income rank from $T-8$ to $T-4$. We assign individuals a pre-period General Practitioner (GP) based on the GP individuals saw the most from $T-3$ to $T-1$.

Separation of Doctors

For the movers and the control sample of non-movers, we measure the separation from their pre-period doctor by the change in the probability that the individual sees their pre-period doctor in the years after the move.

We find that movers have a much larger decrease in the probability they see their pre-period physician after the move than the control sample of non-movers. Figure 2.3 plots the fraction of movers and non-movers that visit their pre-period doctor in each year starting three years prior to the move and going up to to three years after the move. Prior to the move, movers and non-movers see their doctor with a probability equal to approximately 80% in each year. After the placebo-assigned move, the probability non-movers see their physician falls steadily by approximately 5pp a year, so that by three years after the move, the probability non-movers see their pre-period doctor is 60%. Even though non-movers don't actually move, there is a natural fall-off in the probability they see their pre-period physician due to a natural separation rate in

each given year.[29] However for movers, the probability they see their doctor after the move falls sharply to approximately 20% immediately after their move, and decreases to 17% by the third year after the move.[30] This shows that movers have a significantly larger treatment in terms of not seeing their prior doctor after the move.

Computing the Instrumented Change in Prescription Rates

The relevant treatment for our identification strategy is the relative predicted change in physician prescribing rates based on the pre-period physician's prescription drug rate between the movers and the non-mover sample. Denote the prescription rate of doctor j, i's doctor ***B***efore the move: $p_{j(iB)}$, and denote the prescription rate of doctor j', i's doctor ***A***fter the move: $p_{j(iA)}$.[31] We want to estimate: $\mathbb{E}(p^T_{j(iA)} - p^T_{j(iB)}|p^T_{j(iB)}) - \mathbb{E}(p^C_{j(iA)} - p^C_{j(iB)}|p^C_{j(iB)})$ - the difference in conditional expectations based on the pre-period physician's prescribing rate between the treatment and control sample.

To calculate this, we first calculate the change in physician prescription drug rates of their post-period and pre-period doctor for every individual in the mover and non-mover sample. We then plot the relationship between the change in prescribing rates and the pre-period physician's prescribing rate in Figure 2.4A, by binning the pre-period physician's prescribing rate of opioids into 20 equal size bins, and plotting the average change in physician's prescribing rate for each bin. The light red diamonds and fitted line show the relationship between the change in physician's opioid prescribing rate and the pre-period physician's opioid prescribing rate for the movers, while the dark blue dots and line show the relationship for the non-movers.

This figure shows that for movers, the change in physician's opioid drug rate is on average equal to -.75 of their old physician's prescription drug rate. The R^2 for the comparable regression is .39. If there was perfect sorting of individuals to their new physician based on their old physician prescribing rate, we would expect a coefficient of 0, if there was no sorting and perfect mean

[29]Note this does not happen in the pre-period since the assignment to the pre-move doctor is based on the doctor they saw the most in years $T-3, T-2, T-1$.

[30]One reason the probability of seeing the pre-period doctor doesn't fall completely to 0% even three years after the move is that some individuals may move to nearby municipalities and are able to still see their doctor.

[31]We use a leave-out procedure so that $p_{j(iB)}$ and $p_{j(iA)}$ are not calculated using individual i's own prescription drug use.

Figure 2.4: *The Average Change in Physician Prescribing Rate by the Pre-period Physician Prescribing Rate for Movers and Non-Mover Control Group*

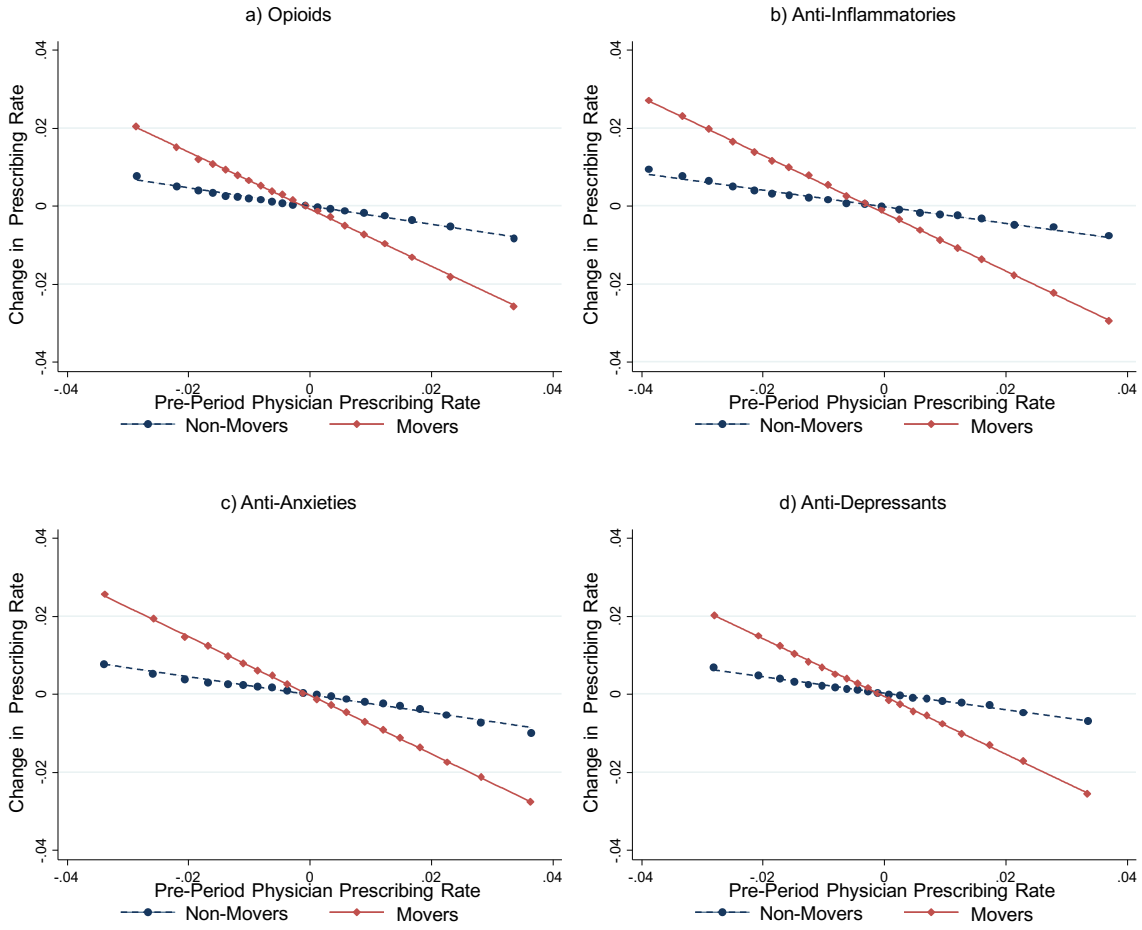

Notes: This figure plots the average change in physician prescribing rates by the pre-period physician prescribing rate for movers (red dots and line) and non-movers (blue dots and line). We bin the pre-period physician prescribing rates into 20 equal sized bins and plot for the x-axis value, the average pre-period physician prescribing rate within that bin, and for the y-axis value, the average change in prescribing rate for individuals within that bin. The x-axis is in terms of rates, while the y-axis is in terms of change of rates. So in the x-axis a .02 physician prescribing rate of opioids means that 2% of the physician's patients take opioids. Panel (a) plots the results for opioid physician prescribing rates. Panel (b) plots the results for anti-inflammatory physician prescribing rates. Panel (c) plots the results for anti-anxiety prescribing rates, and finally Panel (d) plots the results for anti-depressant prescribing rates.

reversion, we would expect a coefficient of -1. The coefficient of -.75 suggests that there is some sorting of individuals into new physicians based on their old physician's prescribing rate, but the mean reversion effect is very strong.

For non-movers, the pre-period's physician drug rate is also predictive of the change in prescription drug rate, but the gradient is much flatter - the change is approximately -.22 of the pre-period physician's prescription drug rate. This is not equal to zero because some of

non-movers separate from their physician and have some mean reversion in their choice of new physicians as well. Therefore, the relative treatment for the movers compared to the non-movers is the difference: $(-.75 - .22) * p_{j(iB)} = -.53 * p_{j(iB)}$. Figure 2.4 Panels B-D show that the results for the other drugs are very similar.

Formally, for each type of prescription drug, we estimate the relative predicted difference of the change in prescription drug rates using the following regression equation:

$$p_{j'(iA)} - p_{j(iB)} = \beta_0 + \beta_1 p_{j(iB)} + \beta_2 Mover_i + \beta_3 p_{j(iB)} \times Mover_i + \epsilon_i \tag{2.2}$$

We calculate the relative predicted difference in the change of prescription drug rates as the product $\hat{\beta}_3$ and the previous physician's prescription drug rate: $\hat{\Delta}_{j(i)} = \hat{\beta}_3 p_{j(iB)}$. Table 2.1 row 3 reports the standard deviation of $\hat{\Delta}_{j(i)}$ for each drug. For opioids, the standard deviation in the relative treatment is .9 percentage points. While it is smaller than the total variation in physician prescription rates, it is still substantial. The variation for the other drugs is slightly higher, with anti-inflammatories having the largest at 1.5 percentage points.

2.4 Effect of Physician Prescribing Rates on Drug Use

This section describes and implements the methodology we use to identify the effect that physician's prescribing rates have on individual's prescription drug use. We find that physicians have a significant causal impact on individual's prescription drug behavior.

2.4.1 Estimating Equation for Impacts of Moves on Prescription Drug Use

To identify the effects of physician prescribing rates on individual's drug use, we use the following identification equation

$$Drug_{it} = \theta_{r(i,t)} \hat{\Delta}_{p(i)} \times Mover_i + \Gamma_{r(i,t)} \hat{\Delta}_{p(i)} + \Psi_{r(i,t)} Mover_i + \mu_{o,d,r} + \epsilon_{it} \tag{2.3}$$

Where $\hat{\Delta}_{p(i)}$ is the relative predicted change in physician prescription drug rates between the mover and non-mover sample based on the pre-period physician's prescribing rate, and $\mu_{o,d,r}$ are origin by destination by year relative to move fixed effects, which control for location based effects. $Drug_{it}$ is an indicator (0 or 1) for whether individual i purchases the drug in year t. $\hat{\Delta}_{p(i)}$ ranges

approximately from -.04 to .04 - such that $\hat{\Delta}_{p(i)} = .02$ would indicate that there was a predicted relative change of 2pp in the physician prescribing rates. $\theta_{r(i,t)}$ is a flexible function allowing for separate coefficients on $\hat{\Delta}_{p(i)}$ for each year relative to the move. We normalize θ_{-1} equal to zero so that the other coefficients indicate the effect of $\hat{\Delta}_{p(i)}$ relative to the year prior to the move. Thus θ_s estimates the triple difference effect - the effect of the predicted change in physician prescription drug rates, for movers relative to non-movers, and in year s relative to the year prior to the move on individual's own prescription drug use. For ease of description, we sometimes refer to θ_s as the effect of a change in physician prescription drug rates in year s relative to the year prior to the move.

The specification assumes that there is a linear effect of the change in physician prescribing rates on individual's prescription drug use. In Appendix B.5, we allow there to be a non-parametric affect of the change in physician prescribing rates on drug use and find that our linear choice is justified.

We include observations for up to three years prior to the move and three years after the move, for the years 1995-2015 and individuals aged 30-70. Note that this is not a balanced sample by individual and year relative to move; however, due to the exact match between the treatment and control group on age and year, each sample is unbalanced in the same way.

2.4.2 Results

We find that an increase in physician prescribing rates of opioids increases individual's own prescription drug use, suggesting that physicians have a causal effect on their patients' use of prescription opioids. Figure 2.5 plots the coefficients ($\theta_{r(i,t)}$) on the instrumented change in physician prescription rates ($\hat{\Delta}_{p(i)}$) for the mover sample, relative to the non-moving sample for each year relative to the year prior to the move. Panel A depicts the results for opioids and shows that in the year of the move, movers' drug use starts increasing with respect to the change in physicians' prescribing rates relative to the response of non-movers. It continues to increase until the year after the move at to a coefficient of approximately .4.[32]

One concern is that time-varying selection biases our results. If this were the case, we would

[32]There is a partial response in the year of the move because on average individuals will spend half of that year in the new municipality.

Figure 2.5: *The Effect of Physician Prescribing Rates on Individual Prescription Drug Use*

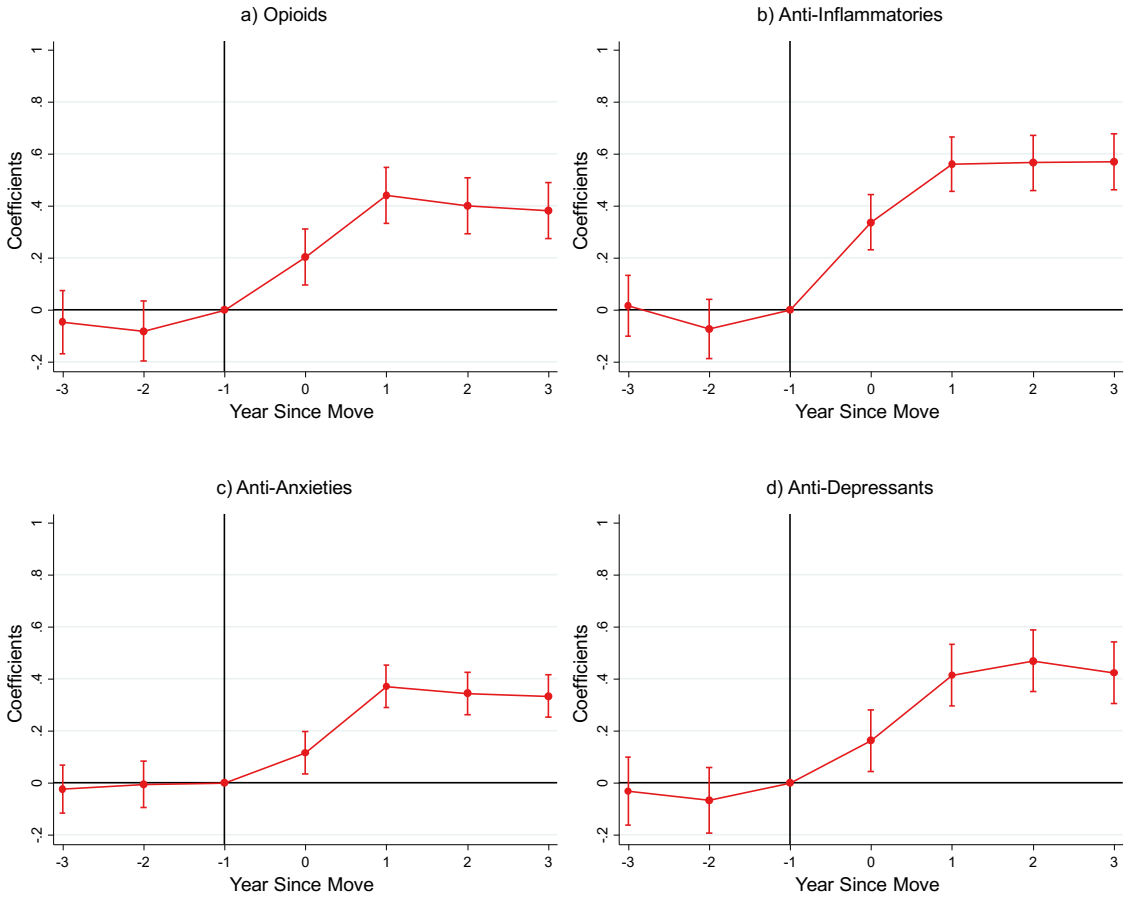

Notes: This plots the coefficients $\theta_{r(it)}$ from the estimate Equation 2.3 by the year since the move. The exact interpretation of the coefficients is the change in the relationship between prescription drug use and the relative instrumented change in physician prescribing rates for the year relative to the year prior to the move $(T-1)$, and for movers relative to non-movers. However given our assumptions, the coefficients can also be interpreted as the effect of a 1 percentage point increase in physician prescribing rates leads to a X percentage point change in drug use. We calculate the relative instrumented change in physician prescribing rates in as the difference in the linear fits from Figure 2.4 of the change in the physician prescribing rates for a particular pre-period physician prescribing rate between the mover and the non-mover sample. Included in the regression are origin by destination by year since event fixed effects (for non-movers the origin and destination are the same). This regression is estimated over the years 1995-2015 for individuals aged 30-70. The bars plot the 95% confidence interval for the coefficients. Panel (a) plots the results for opioid physician prescribing rates. Panel (b) plots the results for anti-inflammatory physician prescribing rates. Panel (c) plots the results for anti-anxiety prescribing rates, and finally Panel (d) plots the results for anti-depressant prescribing rates.

expect to see that the correlation between individual's drug use and the change in physician prescribing rates to increase prior to the move. In fact, there is no trend prior to the move, indicating that the change we see at the time of the move is unlikely to be due to time-varying selection. We may still worry that the relationship between concurrent health *shocks* and the pre-period physician's prescription drug rates is different for the movers than for the non-movers.

Particularly, we would worry that movers who have low-prescribing opioid doctors are more likely to have a concurrent "bad" health shock than individuals who do not move, and movers who have a high-prescribing opioid doctors are more likely to have a "good" health shock than the individuals who do not move.

To check if differential simultaneous shocks are a problem, we look at the relationship of prescription drug use and the change in physician prescribing rates by months since the move. We run the following regression:

$$Drug_{im} = \theta_{r(i,m)}\hat{\Delta}_{p(i)} \times Mover_i + \Gamma_{r(i,m)}\hat{\Delta}_{p(i)} + \Psi_{r(i,m)}Mover_i + \mu_{o,d,r} + \epsilon_{im}, \tag{2.4}$$

where $\theta_{r(i,m)}$ are indicators for months since the move and range from 6 months before the move to 6 months after the move, $Drug_{im}$ is an indicator for whether individual i ever took the prescription drug in month m (rather than within a year), and $\mu_{o,d,r}$ are destination by origin by month relative to move fixed effects. $\hat{\Delta}_{p(i)}$ is the same instrumented change in physician prescription drug rates based on the pre-period physician's prescribing rate.[33]

Figure 2.6 plots $\theta_{r(i,m)}$ by the month since the individual moved for each drug. Panel A plots the coefficients for the outcome of opioid use and the treatment of physician opioid prescriptions. While the coefficients are substantially noisier than the annual coefficients, they also show a flat trend prior to the move, and a distinct increase right at month zero - the month of the move. The fact that the response is immediate at the month of the move provides evidence that this is a causal effect of the move rather than due to correlated concurrent health shocks.

We have shown that these results are unlikely to be driven by differential health trends or shocks. To aggregate the results into one coefficient we estimate:

$$\begin{aligned} Drug_{it} = \theta After_{it} \times \hat{\Delta}_{p(i)} \times Mover_i + \Gamma After \times \hat{\Delta}_{p(i)} + \Psi After_{it} \times Mover_i + \mu_{o,d,r} \\ + \beta_1\hat{\Delta}_{p(i)} + \beta_{r(i,t)}X_{it} \times Mover_i + \epsilon_{it} \end{aligned} \tag{2.5}$$

In Table 2.1, we report θ, which estimates the average effect of the instrumented change in prescription drugs of the movers relative to the non-movers in the three years after the move

[33]Note, however, that the physician prescribing rates are still in annual units: the fraction of their patients that took the prescription drug within the year. This means that the coefficients $\theta_{r(i,m)}$ are not on an equivalent scale as $\theta_{r(i,t)}$.

Figure 2.6: *The Effect of Physician Prescribing Rates on Monthly Individual Prescription Drug Use*

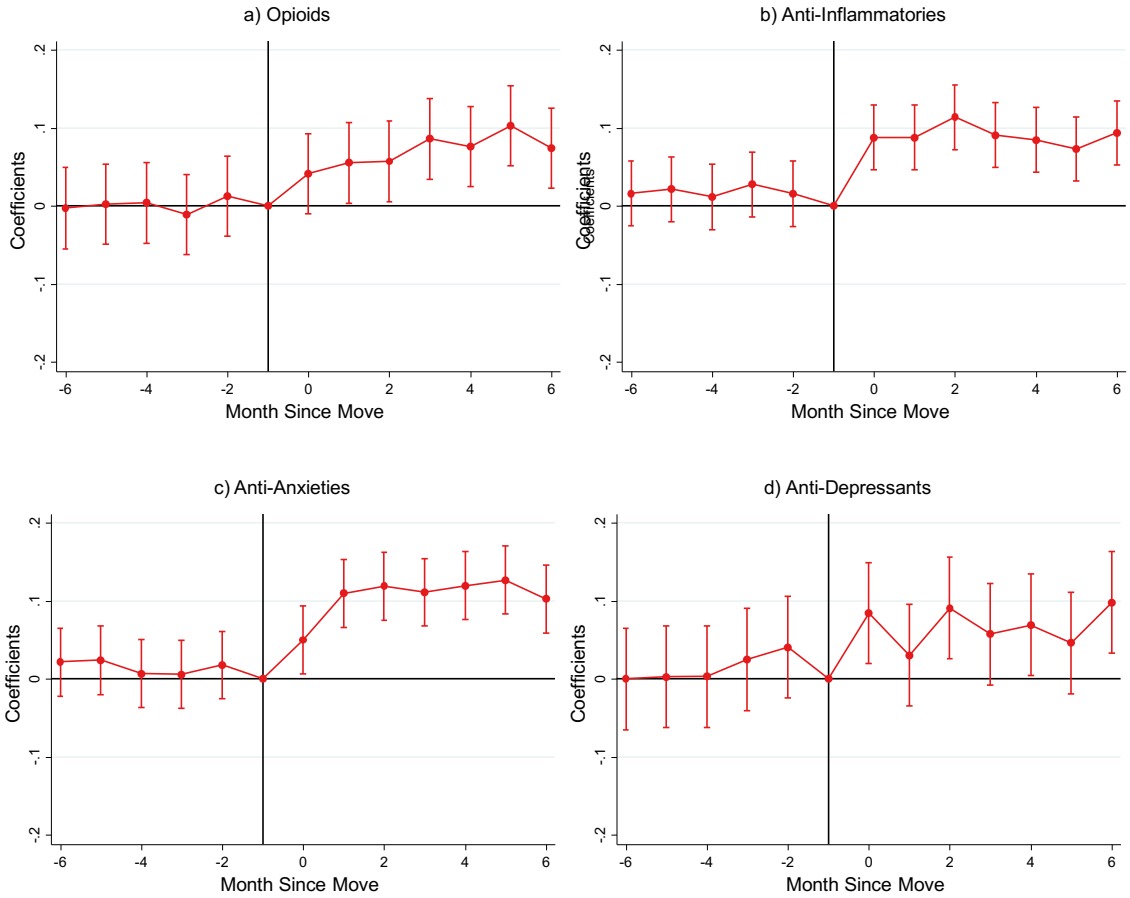

Notes: This figure duplicates Figure 2.5, but instead of annual prescription drug use as the outcome variable, the outcome variable is monthly prescription drug use. Additionally, the relationship between drug use and the relative instrumented change in physician prescribing rates for movers relative to non-movers is calculated for the 6 months prior to the move up until 6 months after the move. Thus this figure plots the coefficients $\theta_{r(im)}$ from the estimate Equation 2.4 by the month since the move. See Figure 2.5 for how we calculate the relative instrumented change in physician prescribing rates. The bars plot the 95% confidence interval for the coefficients. Panel (a) plots the results for opioid physician prescribing rates. Panel (b) plots the results for anti-inflammatory physician prescribing rates. Panel (c) plots the results for anti-anxiety prescribing rates, and finally Panel (d) plots the results for anti-depressant prescribing rates.

versus the three years prior. In column (1) we report the results for opioids and find that an individual's drug use increases by .48pp for a 1 percentage point increase in their physician's prescription drug rate, with a standard error of .03pp.[34] This indicates that 48% of the variation in physician residualized prescription opioid drug rates is driven by the doctor rather than

[34]Another way to interpret the magnitudes coefficient is to put it in terms of standard deviation effects: a one standard deviation increase in physician prescribing rates of opioids leads an individual to increase their own prescription drug use by .8 pp, or 12% of the mean opioid use.

differences in demand of their patients.[35]

Since the physician prescription drug rates are already residualized with respect to individuals age, gender, and education, to calculate the percent of the *total* variation in physician prescribing rates that is due to causal physician effect, we do the following calculation: $\frac{sd_R^2}{sd_{tot}^2} \times .48 = \frac{1.8^2}{2.3^2} \times .48 = .29$. Where sd_R^2 is the squared standard deviation of the residualized physician opioid prescribing rate (Table 2.1 Row 2), and sd_{tot}^2 is the squared standard deviation of the total raw physician prescribing rates (Table 2.1 Row 1). This calculation shows that 29% of the total variation in physician prescribing rates is due to causal physician effects.

Figure 2.5 Panels B-D show the results for the prescribing rates of anti-inflammatories, anti-anxieties, and anti-depressants respectively. They all show a flat trend prior to the move and a strong significant increase after the move. The magnitudes of the increase vary slightly, but they all show a significant response with coefficients between .35 and .6. When we look at monthly prescription drug use in the six months before the move and six months after the move for these drugs in Figure 2.6 Panels B-D, we see that there are no differential pre-trends in the monthly leading up to the move for anti-inflammatoris, anti-anxietes, and anti-depressants either. Additionally, all show an impact on drug use from the change in prescribing rates in the month of the move providing evidence that the effect of the change in physician prescribing rates on individual's prescription drug use is causal.

In Table 2.2 columns 2-4, we report $\hat{\theta}$, from Equation 5, the average effect of the instrumented change in prescription drugs of the movers relative to the non-movers in the three years after the move versus the three years prior. For anti-inflammatories: a 1pp increase in the physician prescribing rate leads to a .58pp increase in individual's probability of taking anti-inflammatories. For anti-anxieties: a 1 pp increase in the anti-anxiety prescribing rate leads to a .36pp increase in the probability of taking anti-anxieties. And for anti-depressants: 1 pp increase in the anti-depressant prescribing rate leads to a .47pp increase in the probability of taking anti-depressants.[36]

[35]This is not equal to 100%, since the residualized physician prescribing rate still includes variation due to differences in demand of their patients based unobservable differences in health.

[36]In terms of standard deviation units: a one standard deviation increase in physician anti-inflammatory prescribing rates leads to a 1.7pp or a 8.8% increase in anti-inflammatory prescription drug use. A one standard deviation increase in physician anti-anxiety prescribing rates leads to a .9pp or a 12% increase in anti-anxiety prescription drug use. A one standard deviation increase in physician anti-depressant prescribing rates leads to a .8pp or a 10% increase in prescription anti-depressant use.

Table 2.2: *The Effects of Physician Prescribing Rates on Prescription Drugs*

	Prescription Drug Use			
	Opioids (1)	Anti-Inflammatories (2)	Anti-Anxieties (3)	Anti-Depressants (4)
After × Mover × Predicted Δ in Physician Prescribing Rates	0.448***	0.585***	0.359***	0.467***
	(0.0360)	(0.0338)	(0.0273)	(0.0401)
N	15,324,329	15,324,329	15,324,329	15,324,329

Notes: This table presents the estimated coefficients from Equation (2.5) in the paper: a regression of individual prescription drug use (an indicator variable) on the interaction of an indicator for after the move, whether the individual is a mover, and the relative predicted change in physician prescribing rates based on the pre-period physician's prescribing rate. Also in the regression includes the full set of pairs of interactions between those three variables, as well as origin by destination by an indicator for after fixed effects. The regression includes individuals who move across municipalities at time T=0, as well as a matched set of individuals who do not move who are matched by year, age, gender, education, quartiles of prior physician prescribing rates, and a quartile for the individual's average income rank from T-8 to T-4. Observations up to three years after the move and three years prior to the move are included in the regression, while the year of the move is not. Individuals who are aged 30-70 and from the years 1995-2015 are included. The relative predicted change in physician prescribing rates is the difference for the mover and non-mover sample of the linear prediction from a regression of change in physician prescribing rates from after the "move" from before the "move" on their prior physician prescribing rates. Column (1) presents the results when the relative predicted change in prescribing rates is for opioids, and the outcome measure is opioid prescription drug use. Column (2) presents the results when the prescribing rate and the outcome measure is for anti-inflammatories. Column (3) presents the results when the prescribing rate and the outcome measure is for anti-anxieties, and Column (4) presents the results for when the outcome measure is for anti-depressants. Standard errors are reported before the coefficients and calculated by clustering at the individual level. Given that the dependent variable is an indicator, and the treatment variable is in terms of a rate, the interpretation of coefficient in column (1) is the effect of a 1 percentage point increase in the relative predicted change in physician prescribing rates leads to a .45 percentage point increase in opioid prescription drug use for movers relative to non-movers, after the move compared to before the move.

2.4.3 Heterogeneity In Treatment Impacts by Individual Characteristics

To understand the distributional implications of the variation in physician prescribing rates, we need to know which individuals are the most influenced by their physicians. To do so, we estimate heterogenous effects of physician prescribing rates on prescription drug use by gender, age, education, and whether individuals work in blue collar occupation. We find a substantial amount of variation in the effects by these characteristics especially for prescription opioid use.

We turn age and education into binary variables based on whether the individual's level of a given variable is above or below the median at the time of the move.[37] We create an indicator for whether an individual worked in a blue collar occupation based on their occupation four years prior to the move.

To estimate heterogenous effects by each characteristic, X_i, we first reestimate $\hat{\Delta}_{p(i)}$, the relative predicted change in physician prescribing rate based on their pre-period physician's prescription rate for movers compared to non-movers, as a function of X_i. We do this because individuals with different characteristics may have different amounts of mean reversion from the pre-period to their post-period physician prescribing rate.

Since each X_i is binary, we simply estimate:

$$p_{j(iA)} - p_{j(iB)} = \beta_0^0 + \beta_1^0 p_{j(iB)} + \beta_2^0 Mover_i + \beta_3^0 p_{j(iB)} \times Mover_i + \epsilon_i \quad \forall i \text{ s.t. } X_i = 0 \tag{2.6}$$

$$p_{j(iA)} - p_{j(iB)} = \beta_0^1 + \beta_1^1 p_{j(iB)} + \beta_2^1 Mover_i + \beta_3^1 p_{j(iB)} \times Mover_i + \epsilon_i \quad \forall i \text{ s.t. } X_i = 1 \tag{2.7}$$

Where $\delta_{j(i)}$ is the change in physician prescribing rates: $p_{j(iA)} - p_{j(iB)}$). We then create the new instrumented variable as a function of X_i: $\hat{\Delta}_{p(i)i} = \beta_3^0(1 - X_i) \times p_{j(iB)} + \beta_3^1 X_i \times p_{j(iB)}$, and estimate the following equation to identify heterogeneous effect of physician's prescribing rates on individual use:

$$\begin{aligned} Drug_{it} =& \theta_1 After \times \hat{\Delta}_{p(i)i} \times Mover_i \times X_i + \Gamma_1 After \times \hat{\Delta}_{p(i)i} \times X_i + \\ & \Psi_1 After \times Mover \times X_i + \beta_1 Mover \times X_i + \lambda_1 \hat{\Delta}_{p(i)i} \times X_i + \\ & \theta_0 After \times \hat{\Delta}_{p(i)i} \times Mover_i + \Gamma_0 After \times \hat{\Delta}_{p(i)i} + \\ & \Psi_0 After \times Mover + \beta_0 Mover + \lambda_0 \hat{\Delta}_{p(i)i} + \mu_{o,d,r,X_i} + \epsilon_{it} \end{aligned} \tag{2.8}$$

[37] The median age is 42 and the median education level is 14 years of education.

θ_1 estimates the difference between the effect of a change in physician prescription drug rates for individuals with characteristic X_i compared to those without X_i, while θ_0 estimates the base effect of the change in physician prescription drug rates for those without characteristic X_i. We also include origin by destination by after by characteristics fixed effects (μ_{o,d,r,X_i}) to control for the possibility that individuals with different X_i are differentially influenced by location effects as well. Table 2.3 reports the estimates for $\hat{\theta}_1$ and $\hat{\theta}_0$ with Column 1 showing heterogeneity by age ($X_i = age_i > 42$), Column 2 by gender ($X_i = female_i$), Column 3 by education ($X_i = YearsEduc_i > 14$), and Column 4 by occupation ($X_i = BlueCollar_i$).[38]

We find that individuals who are older, female, less educated, or worked in blue collar occupations have a significantly larger effect of a change in their physician's opioid prescribing rate on their own opioid prescription drug use. Table 2.3 Panel A reports the coefficients for opioids, and shows that older individuals have a 75% larger response than younger individuals; women have a 60% larger response than men; less educated individuals have a 75% larger response than more educated individuals; and blue collared workers have a 89% larger response than white collared workers. Tables 2.3 Panels B-D show the results for anti-inflammatories, anti-anxieties, and anti-depressants respectively. For anti-inflammatories, we find that those with less education are more affected by their physician's prescribing rate. For anti-anxieties in Table 2.3 Panel C, we see that older workers, women, and blue collar workers are more likely to be affected by their physician prescribing rate. For anti-depressants in Table 2.3 Panel D, we find that women and blue collar workers have a larger effect.

Generally, we find that differences in the effects align with the differences in the exante probability individuals use the drug. Appendix Table B.2, Panels A-D report the difference in probabilities of prescription drug use for individuals with these different characteristics for the different prescription drugs we study. Those who are more likely to use the drug are also more likely to be affected by the change in physician prescribing rates.

[38]Where 42 is the median age, and 14 is the median years of education in the sample.

Table 2.3: *Heterogeneity of Treatment Effects on Prescription Drugs*

	Prescription Drug Use			
X:	Old	Female	Educated	Blue-Collar
	(1)	(2)	(3)	(4)
A. Opioids				
After×Δ×Mover	0.342***	0.353***	0.533***	0.279***
	(0.0464)	(0.0464)	(0.0428)	(0.0474)
After×Δ×Mover×X	0.227***	0.211**	-0.228***	0.249***
	(0.0663)	(0.0666)	(0.0676)	-0.0692
B. Anti-Inflammatories				
After×Δ×Mover	0.548***	0.549***	0.652***	0.554***
	(0.0458)	(0.0460)	(0.0419)	(0.0484)
After×Δ×Mover×X	0.0695	0.0968	-0.168*	0.0931
	(0.0647)	(0.0649)	(0.0658)	(0.0716)
C. Anti-Anxieties				
After×Δ×Mover	0.305***	0.318***	0.384***	0.274***
	(0.0345)	(0.0350)	(0.0326)	(0.0350)
After×Δ×Mover ×X	0.132**	0.0987*	-0.0734	0.125*
	(0.0498)	(0.0498)	(0.0506)	(0.0524)
D. Anti-Depressants				
After×Δ×Mover	0.471***	0.386***	0.520***	0.404***
	(0.0505)	(0.0514)	(0.0475)	(0.0517)
After×Δ×Mover×X	0.0141	0.154*	-0.134	0.161*
	(0.0727)	(0.0725)	(0.0736)	(0.0769)
N	15,324,329	15,324,329	15,324,329	12,407,329

Notes: This table presents the estimated coefficients from Equation 2.8 in the paper: a regression of prescription drug use on the full interaction between an indicator for after the move, whether the individual is a mover, the predicted relative change in the physician prescribing rate, and an indicator for characteristic X. We report the coefficients of the effect of physician prescribing rates for those with out the characteristic (e.g. the first row of Panel A), and the difference in the effect for those with the characteristic versus those without the characteristic (e.g. the second row of Panel A). Panel A reports the results for opioids, while Panel B is for anti-inflammatories, Panel C is for anti-anxieties, and Panel D is for anti-depressant prescribing rates and anti-depressant prescription drug use. Column (1) presents the results when the indicator is for individual being older than 42 at the time of the move, while Column (2) is for an indicator of female, Column (3) for an indicator of having more than 14 years of education, and column 4 for individual being in a blue collar occupation in T-4. It includes the sample sample as Table 2.2.

2.5 Effect of Physician Prescribing Rates on Labor Outcomes

In this section, we estimate the effect of physician prescribing rates on individual's labor supply outcomes. We first present the main results, which look at the effects on labor income rank. Next, we show the effects on additional labor outcomes: labor force participation, log labor income, labor income rank defined for individuals with positive labor income, two sick pay measures, and Disability Insurance receipt. Finally, we show that individuals who move farther distances and have a larger change in prescribing rates also have bigger changes in labor supply. This provides additional evidence that the labor supply effects come from the change in the physician's prescribing rate rather than other confounding factors. In Appendix B.6, we also estimate whether there are heterogenous effect for individuals by their age, gender, education, and occupation; however, due to large standard errors, it is difficult to make any strong conclusions.

2.5.1 Estimating Equation for Effects of Physician Prescribing Rates on Labor Income

Similar to our estimation of the effects of physician prescription rates on drug utilization, our equation for identifying the effects of a change in physician prescribing rates on labor income is the following:

$$LaborIncome_{it} = \theta^{L}_{r(i,t)}\hat{\Delta}_{p(i)} \times Mover_i + \Gamma^{L}_{r(i,t)}\hat{\Delta}_{p(i)} + \Psi^{L}_{r(i,t)}Mover + \mu^{L}_{o,d,r} + \epsilon_{it} \tag{2.9}$$

Where $\hat{\Delta}_{p(i)}$ is the predicted relative change in physician prescription drug rates based on the pre-period physician's prescribing rate (calculated in Section 3.3.2), and $\mu^{L}_{o,d,r}$ are origin by destination by year relative to move fixed effects, which control flexibly for local labor market effects of the origin and destination. $\theta^{L}_{r(i,t)}$ is a flexible function allowing for separate coefficients on $\hat{\Delta}_{p(i)}$ for each year relative to the move. We normalize θ_{-1} equal to zero so that the other coefficients indicate the effect of $\hat{\Delta}_{p(i)}$ relative to the year prior to the move. Thus, just like in section 4.1, θ^{L}_{s} estimates the triple difference effect - the effect of the predicted change in physician prescription drug rates, for movers relative to non-movers, and in year s relative to the year prior to the move on individual's labor income.

The primary outcome variable we consider is labor income rank because it includes both the extensive and intensive margin responses. In section 5.2.2, we look at the effect of physician

prescribing rates on other labor market outcomes to see if the effects are robust to different measures and to identify whether the results are driven primarily by the extensive or intensive margin response.

2.5.2 Results

Results on Labor Income Rank

Opioids: We find that an increase in physician opioid prescribing rates leads to a decrease in individual's labor income. Figure 2.7 plots the coefficients ($\theta^{L}_{r(i,t)}$) on the instrumented change in physician prescription rates ($\hat{\Delta}_{p(i)}$) for the moving sample relative to the non-moving sample for each year relative to the year prior to the move. Panel A shows the results for the effect of opioid prescribing rates on labor income rank. It shows that in the year of the move, mover's labor income rank starts to decrease with respect to the change in the physician's opioid prescribing rate relative to the response of non-movers. It continues to decrease until the year after the move at which point it levels off. Prior to the move, there is no significant trend, which means that individuals who move do not change their labor supply in response to the change in their physician's opioid prescribing rate prior to the move relative to the non-mover sample. This suggests that the large change we see at the time of the move is unlikely to be due to time-varying selection. Table 2.4 Panel A reports the estimates for the effect of the predicted change in physician opioid prescription rates in the three years after the move versus the three years prior to the move, clustering the standard errors at the individual level. Column (1) reports that a 1pp increase in the opioid physician prescription rate is associated with a .11 percentile decrease in labor income rank.

Because we might be concerned that individuals may move for different reasons and have different effects of the move, Figure 2.7 Panel B show that the results do not change when we include controls to allow for individuals with different pre-characteristics to have different effects of the move. Specifically, we include year since the move fixed effects interacted with an indicator for mover, and separately interacted with the following: a quadratic in age, gender, and the full set of interactions of years of education, average labor income rank from $T-8$ to $T-4$, and age. We find that when we add these controls, there is little change in the results. Table 2.4 Column (2)

Table 2.4: *The Effects of Physician Prescribing Rates on Labor Income Rank*

	Labor Income Rank		
	(1)	(2)	(3)
A. Opioids			
After×Δ×Mover	-0.113***	-0.117***	-0.123*
	(0.0366)	(0.0323)	(0.0467)
B. Anti-Inflammatories			
After×Δ×Mover	0.0334	-0.0354	-0.00156
	(0.0225)	(0.0199)	(0.0223)
C. Anti-Anxieties			
After×Δ×Mover	-0.0878**	-0.0346	0.0546
	(0.0283)	(0.0248)	(0.0345)
D. Anti-Depressants			
After×Δ×Mover	-0.145***	-0.110***	-0.0774*
	(0.0361)	(0.0320)	(0.0382)
Controls			
Ind Pre-Chars x Yr Since Event x Moved		x	x
Other Prescribing Rates			x
N	15324329	15171445	15171445

Notes: This table presents the estimated coefficients from a regression of labor income rank on the interaction of an indicator for after the move, whether the individual is a mover, and the relative predicted change in physician prescribing rates based on the pre-period physician's prescribing rate. Also in the regression includes the full set of pairs of interactions between those three variables, as well as origin by destination by an indicator for after fixed effects. We measure labor income as taxable labor and self-employed earnings. We convert labor income into percentile ranks within an individual's year of birth, the year, and their gender using the full sample of the Danish population (not just the movers and non-mover control sample). It is on a scale from 0 (lowest income) to 1 (highest income). Column (1) replicates the specifications from Table 2.2, substituting labor income rank as the left-hand side variable instead of prescription drug use. Column (2) additionally includes control for individual pre-characteristics that are fully interacted with an indicator for after and an indicator for being a mover. The pre-characteristics include a quadratic in age at the time of move, an indicator for female, and the full interaction of average labor income rank over T-8 to T-4, years of education, and age. Column (3) additionally includes the full interactions of the other prescription rates and indicators for after and being a mover. Thus Column (3) represents coefficients from all the same regression, whereas Column (1)-(2) each report coefficients from four different regressions.

Figure 2.7: *The Effect of Opioid Physician Prescribing Rates on Labor Income Rank*

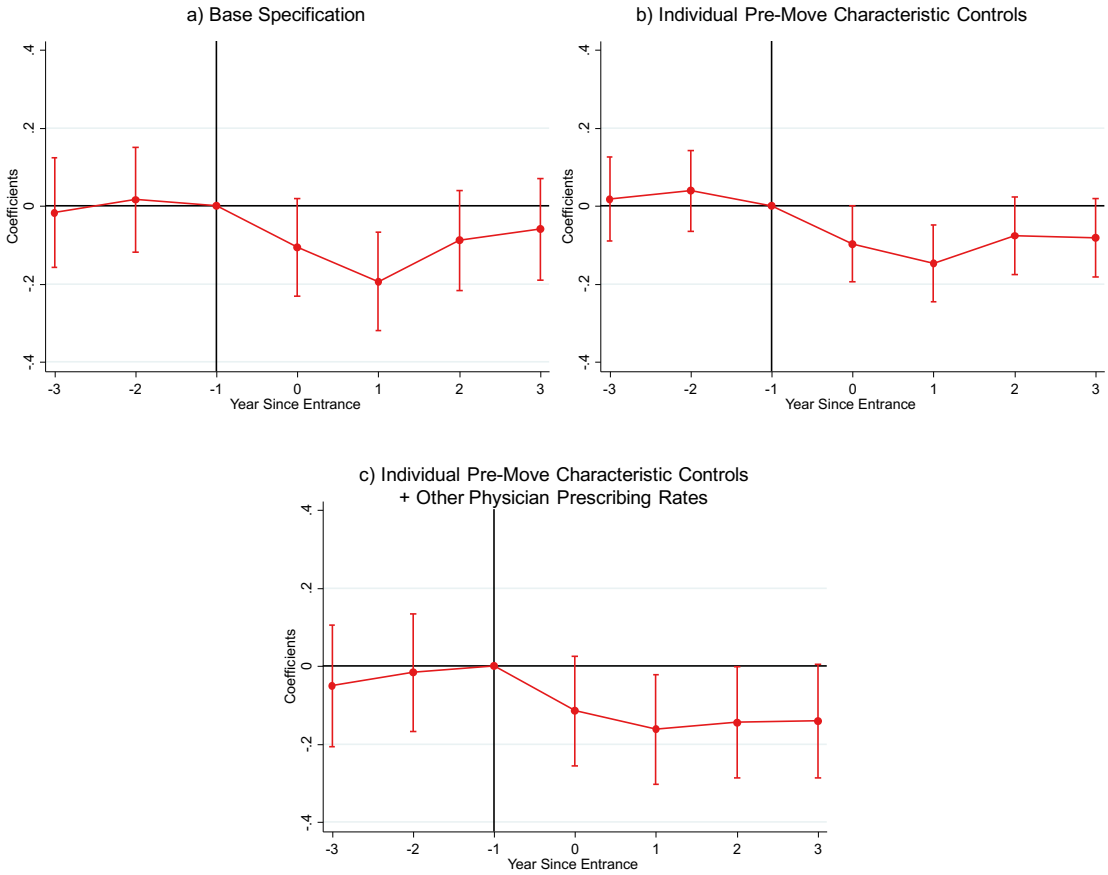

Notes: This figure replicates Figure 2.5a, but instead of annual drug use as the outcome variable, the outcome variable is individual's labor income rank. Specifically, it plots the coefficients $\theta^{L}_{r(it)}$ in Equation 2.9 by the year since the move. We can therefore interpret the coefficients as the change in the relationship between labor income rank and the relative instrumented change in physician opioid prescribing rates for the year relative to the year prior to the move $(T-1)$, and for movers relative to non-movers. If our identifying assumptions hold, we can interpret the coefficients as the effect that a 1 percentage point change in physician prescribing rates has on individual's income in the year since the change. Labor income is defined as the sum of labor income as well as self employment income. We calculate ranks of labor income from the full population of individuals in Denmark and calculate ranks within year, age, and gender groups. It includes individuals with zero labor income and is on a scale of 0 to 1. The regression that the coefficients are from is calculated on the sample of individuals aged 30-70 who move and the non-moving control group from 1995-2013. Panel (a) plots the coefficients for a specification with the same set of controls as Figure 2.5 (a) (municipality origin by destination by year since move fixed effects). Panel (b) includes controls for individual pre-characteristics by year since the move fixed effects and whether the individual is treated. Specifically, we include a quadratic in age at the time of the move, gender, and the full interaction of the average labor income rank from $T-8$ to $T-4$, individual's years of education, and their age at the time of the move. Panel (c) includes the same controls as Panel (b), but additionally includes interactions between the relative instrumented change in physician prescribing rates for the other drugs, the year since the move, and an indicator for whether the individual is treated. Therefore, it controls for the effects of the other physician prescribing rates as well.

reports the aggregated effect of a 1pp change in physician prescribing rates after the move versus before with these controls as a .12 percentile decrease in labor income rank, which is statistically significant at the 5% level.

Figure 2.8: *The Effect of Anti-Inflammatory Physician Prescribing Rates on Labor Income Rank*

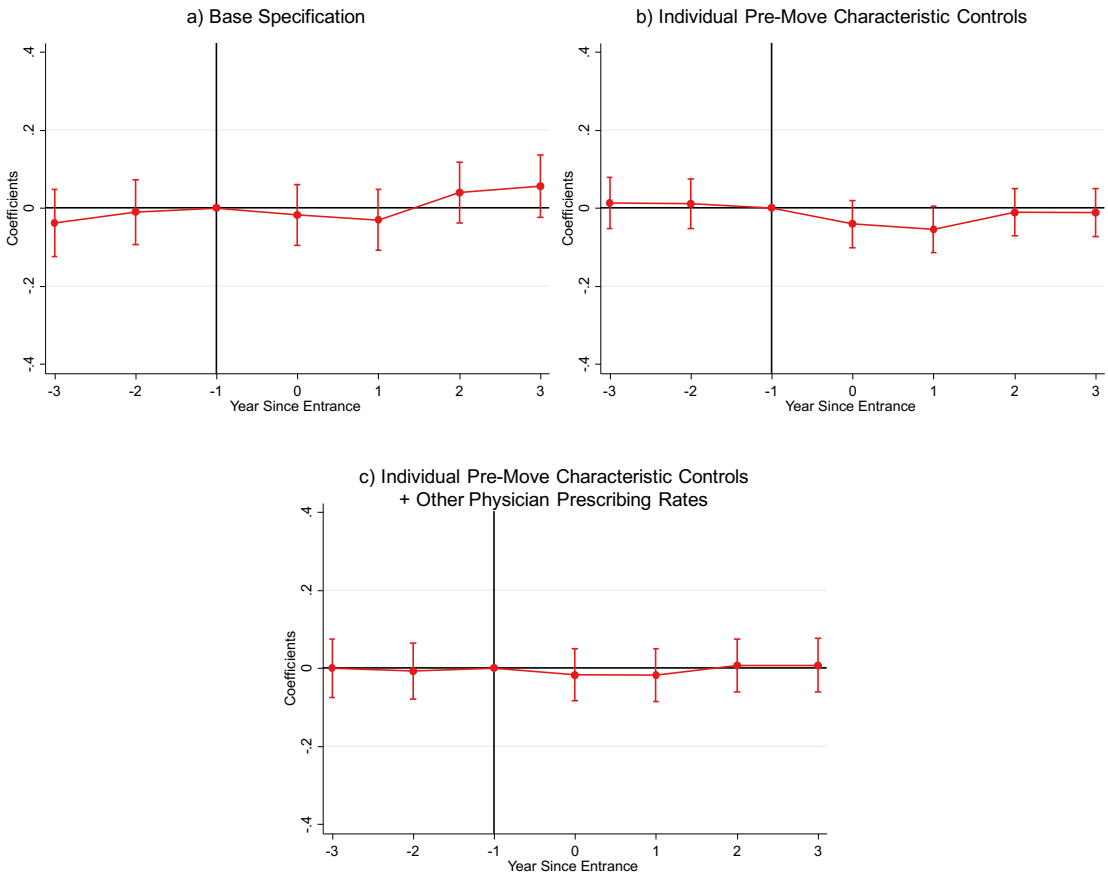

Notes: This figure replicates Figure 2.7, but instead of the intensity of treatment being the predicted change in opioid prescribing rates, here it is the predicted change in anti-inflammatory prescribing rates. See the notes for Figure 2.7 for details.

Anti-Inflammatories: Figure 2.8, Panel A plots the results for anti-inflammatory prescribing rates, and Panel B plots the coefficients for the regression that includes controls allowing for heterogenous effects of the move for individuals with different observable pre-characteristics. We find no significant effects for a change in physician anti-inflammatory prescribing rates on labor income rank in either specification. There is maybe a small decrease in the year after the move, but it is not quite statistically significant at the 5% level. Table 2.4 Panel B, Column (1) and Column (2) report the estimates for the aggregated effect after versus before and finds coefficients of 0.033, with a standard error of .023 and -.035, with a standard error of .02, neither of which are statistically significant at the 5% level.

Anti-Anxieties: For anti-anxieties (Figure 2.9 Panels A and B), we find no discernible effect

Figure 2.9: *The Effect of Anti-Anxiety Physician Prescribing Rates on Labor Income Rank*

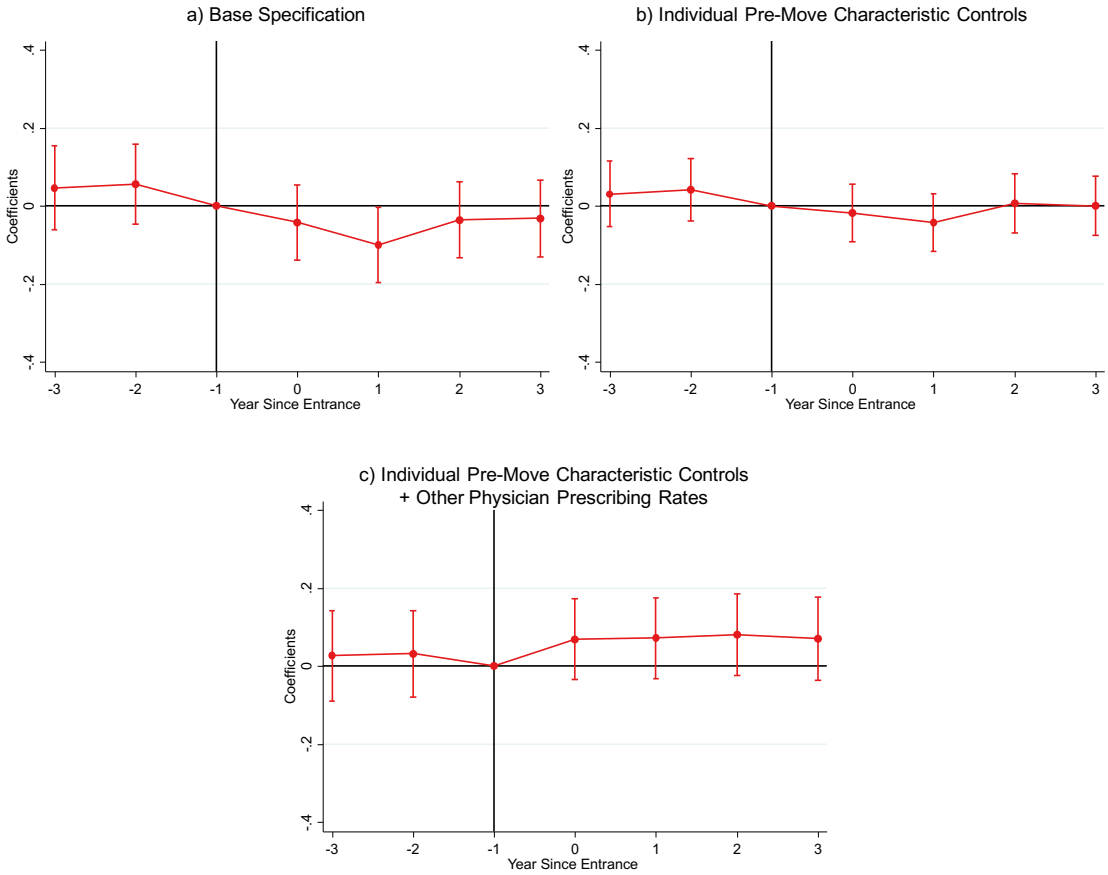

Notes: This figure replicates Figure 2.7, but instead of the intensity of treatment being the predicted change in opioid prescribing rates, here it is the predicted change in anti-anxiety prescribing rates. See the notes for Figure 2.7 for details.

of physician anti-anxiety prescribing rates on labor income rank with or without controls for individual pre-characteristics. Again, there is perhaps a small decrease in the year after the move that is just barely significant at the 5% level, but when we control for differences in individual pre-characteristics, this effect attenuates and becomes statistically insignificant. Table 2.4 Panel C, Column (1) reports the coefficient for the aggregated effect of after versus before as -.088 with a standard error of .028 without controls. Once we include controls for heterogeneity in the effects of the move by individual characteristics in Column (2), this coefficient becomes -.035 with a standard error of 0.02 and is no longer statistically significant at the 5% level.

Anti-Depressants: In Figure 2.10 Panels A and B, we plot the effects of physician anti-depressant prescribing rates on labor income rank, with and without controls for previous individual characteristics. They show that an increase in the physician prescribing rate of anti-depressants

Figure 2.10: *The Effect of Anti-Depressant Physician Prescribing Rates on Labor Income Rank*

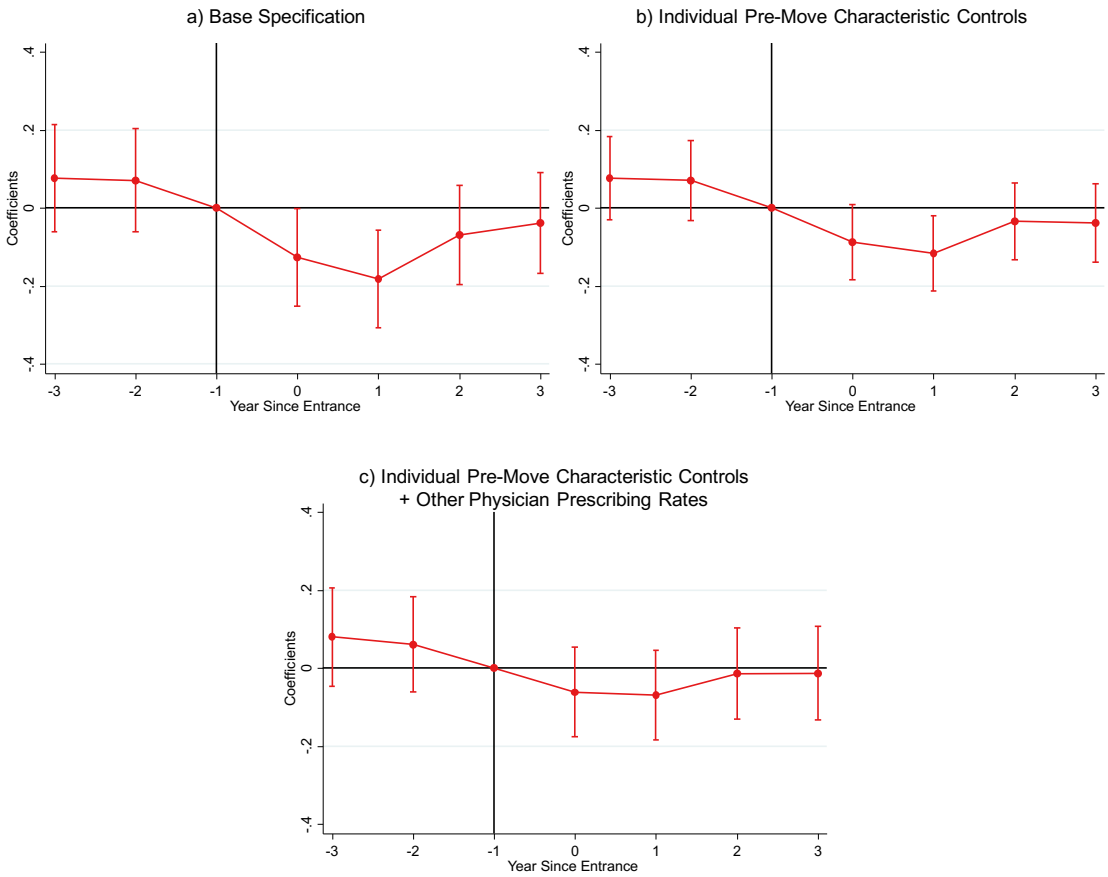

Notes: This figure replicates Figure 2.7, but instead of the intensity of treatment being the predicted change in opioid prescribing rates, here it is the predicted change in anti-depressant prescribing rates. See the notes for Figure 2.7 for details.

leads to a decrease in labor income rank. While not significant, the pre-trends are not entirely flat, which may suggest that some of the effect is due to differential trends for movers and non-movers with respect to the predicted change in the physician prescribing rate of anti-depressants. Aggregating the coefficients from after and before in Table 2.4, we find that without controls a 1 pp increase physician prescribing rates of anti-depressants is associated with a decrease of .15 percentiles in individual's labor income income rank, with a standard error of .036. When we add controls for individual characteristics in Column (2) and Figure 2.10 B, the effect decreases in magnitude to .11, but stays statistically significant at the 5% level. However, this difference includes changes in the pre-period so it is unlikely to reflect a causal effect.

Horse-Race: In the above analysis, we looked separately at the effect of physician prescribing

rates of different drugs. However, physician's prescribing rates for different drugs are highly correlated with each other. Table 2.5 reports the correlation between the different physician prescribing rates, as well as the correlation between the specific treatments we use in our analysis - the relative predicted change in physician prescribing rates. The correlations range between .23 and .57. The highest correlations are between physician opioid prescribing rates and the prescribing rates of the other drugs. For example, for the correlation between th predicted relative change in physician prescribing rates of opioids and anti-depressants is .57, while for opioids and anti-anxieties it is .48, and for opioids and anti-inflammatories it is .43. Because the treatments are correlated, it is unclear whether the estimated effects on labor income is due to the specific prescribing rate or due to it's correlation with the other prescribing rates To separate out the specific effects of each physician prescribing rate, we control simultaneously for them in the same regression. The interpretation of these coefficients, for example, is the effect of having a physician that has a higher opioid prescribing rate holding fixed their prescribing rate of anti-inflammatories, anti-anxieties, and anti-depressants.

Table 2.5: *Correlation Between Physician Prescribing Rates*

Physician Prescribing Rates:	Opioids (1)	Anti-Inflammatories (2)	Anti-Anxieities (3)
A. Correlation of Physician Averages			
Anti-Inflammatories	0.403		
Anti-Anxieties	0.553	0.230	
Anti-Depressants	0.455	0.304	0.362
B. Correlation of Relative Predicted Change			
Anti-Inflammatories	0.428		
Anti-Anxieties	0.476	0.307	
Anti-Depressants	0.565	0.235	0.382

Notes: This table presents the correlations between physician prescription rates of different drugs (Panel A) and the correlation between the relative predicted change in physician prescribing rates. Both are weighted by the estimating sample we us in Table 2.2. Physician prescribing rates are calculated by first taking out effects from immutable individual characteristics, specifically: of age, gender, education and the year. The relative predicted change is estimated as explained in the Notes of Table 2.2, using the difference in the predicted change in prescribing rates based on the pre-move physician for movers and non-movers.

Figure 2.7 Panel C, shows that for opioids there is similar sized effects when we additionally

control for the other prescribing rates of drugs. On the other hand, Figure 2.8 Panel C shows that any small decrease in labor income that was associated with an increase in inflammatory prescribing rates is gone once we include controls for the other prescribing rates. Figure 2.9 Panel C shows the results for anti-anxiety prescribing rates and shows that once we control for the other prescribing rates, an increase in the anti-anxiety prescribing rate is associated with a small *increase* in labor income - though none of the point estimates are statistically significantly different from zero at the 5% level. Finally, Figure 2.10 Panel C shows that once we control for the other prescribing rates, the effect of increase in the physician prescribing rates of anti-depressants on labor income is no longer significant.

In Table 2.4 Column 3, we report the estimates of the aggregate effect of the relative instrumented change after versus before the move when we control for the other prescribing rates for each type of drug. We find that a 1 pp increase in opioid prescribing rate leads to a -.12 (se=.05) percentile decrease in labor income rank; a 1 pp increase in the anti-inflammatory prescribing rate leads to a -.002 (se=.022) percentile change in labor income rank; a 1 pp increase in the anti-anxiety prescribing rate leads to a .055 (se=.035) percentile change in the labor income rank, and a 1pp increase in the anti-depressant prescribing rate leads to a -.077 (se=.038) percentile change in labor income rank.[39] Note that for anti-depressants, that the coefficient is significant, but this is again likely due to the pre-trend we see in Figure 2.10 Panel C. Therefore, this estimated effect is unlikely to be causal.

We conduct a test to see if the opioid effect is statistically significantly different than the other effects. We find that the p-value for opioids and anti-inflammatories to have the same effect is .03, for opioids and anti-anxieties to have the same effect is .003, and for opioids and anti-depressants to have the same effect the p-value is .0008. This provides evidence that the prescribing rates of opioids have a negative effect on labor income rank, while the other drugs have smaller or no effects.

[39]To compare the effect that physician prescribing rates have on labor income once we control for all the prescription drug rates simultaneously to the effect they have on prescription drug use, we rerun the regressions of prescription drug use for each drug on the full set of physician prescribing rates. These results are reported in Appendix Table B.3. Column 1 shows that the coefficient on the opioid physician prescribing rate in the regression on opioid drug use is .61.

Other Labor Outcomes

To understand the effect that physician prescribing rates have on labor supply more completely, we look at the effects on other labor supply measures: labor force participation, log labor income, labor income rank for individuals with positive labor income, two measures of receipt of sick pay, and receipt of disability insurance. We find that physician's prescribing rate of opioids has a negative and significantly significant impacts on individual's labor force participation and their log labor income, while the prescribing rates for the other drugs have no significant effects on any of the additional outcomes we look at.

Table 2.6 shows the results of the effects of physician opioid prescribing rates when these different measures are the outcome variables. We include observations from three years prior to the move and three years after the move, not including the year of the move, and estimate the following equation, where Y_{it}, is the outcome variable of interest:

$$Y_{it} = \theta After_{it} \times \hat{\Delta}_{p(i)} \times Mover_i + \Gamma After \times \hat{\Delta}_{p(i)} + \Psi After_{it} \times Mover_i + \mu_{o,d,r} + \beta_1 \hat{\Delta}_{p(i)} + \beta_{r(i,t)} X_{it} \times Mover_i + \epsilon_{it} \quad (2.10)$$

We report the coefficient, θ, on the interaction with an indicator for $After$, the predicted relative change in physician prescribing rates, $\hat{\Delta}_{p(i)}$, and the indicator for $Mover$.

In Table 2.6, we include the same set of controls as Table 2.4 Column 3, controlling flexibly for differences between individuals and for the physician prescribing rates of other drugs. In Appendix Table B.4, we show the results for when we include controls for differences between individuals, but not other prescribing rates. For each column, we run equation (10) on a different outcome variable. Column (1) reports the results for labor income rank (replication of Table 2.4 Column 3), Column (2) reports the results on labor force participation (either equal to 0 or 1), Column (3) reports the results on $ln(LaborIncome + 1)$, Column (4) reports the results on labor income rank defined for individuals with positive labor income (ranges from 0 to 1), Column (5) reports the results on an indicator for receiving sick pay within a year (either 0 or 1), Column (6) reports the results on an indicator for receiving sick pay for more than four weeks (either 0 or 1), and Column (7) reports the results on DI receipt (either 0 or 1).

Panel A shows the results for the effect of physician prescribing rates of opioids on various

Table 2.6: *The Effects of Physician Prescribing Rates on Other Labor Outcomes with Full Controls*

	Labor Income Rank (1)	LFP (2)	Log Labor Income (3)	Labor Inc. Rank Pos. LFP (4)	Any Sick Pay (5)	Sick Pay > 4 weeks (6)	DI Receipt (7)
A. Opioids							
After×Δ×Mover	-0.123*	-0.196**	-2.316**	-0.0651	0.0481	0.146	0.0109
	(0.0467)	(0.0699)	(0.844)	(0.0476)	(0.0808)	(0.0824)	(0.0470)
B. Anti-Inflammatories							
After×Δ×Mover	-0.00156	-0.0108	-0.182	0.0161	-0.0413	0.0200	0.0405
	(0.0223)	(0.0330)	(0.399)	(0.0225)	(0.0384)	(0.0388)	(0.0220)
C. Anti-Anxieties							
After×Δ×Mover	0.0546	0.0859	1.058	0.0383	-0.0152	-0.0342	-0.0532
	(0.0345)	(0.0512)	(0.619)	(0.0350)	(0.0585)	(0.0596)	(0.0339)
D. Anti-Depressants							
After×Δ×Mover	-0.0774*	-0.0312	-0.773	-0.0615	-0.0487	-0.111	0.0512
	(0.0382)	(0.0573)	(0.691)	(0.0391)	(0.0660)	(0.0680)	(0.0399)
N	15,171,445	15,171,445	15,171,445	12,469,064	8,964,761	8,587,148	13,750,107

Notes: Column (1) of this table replicates Column (3) of Table (2.4). The other columns use the same set of controls but change the outcome variables. The outcome variable for Column (2) is labor force participation which is an indicator variable for having positive labor income. For Column (3), log labor income is defined as ln(Labor Income +1), where labor income is first put into 2015 Danish Kroner (6.5 Dkr 1$) using the Danish CPI. Column (4) replicates column (1) but only includes individuals with positive labor income. Column (5)'s outcome measure is an indicator for any sick pay which is an indicator for whether their employer paid them any sick pay within the year. It is defined only for private sector employees. The outcome variable in Column (6) is an indicator for taking sick pay longer than four weeks, which is identified by whether the municipality paid sick pay for private section employees. Because it also includes payments for maternity leave, we set it equal to missing for women who had a baby within the year or the year previous. Column (7) reports the results for when the outcome variable is Disability Insurance receipt which is defined for individuals less than 65 and is also an indicator variable.

labor supply measures. We find that there is large negative and statistically significant effects on labor force participation and log labor income: a 1 percentage point increase in physician's opioid prescribing rate leads to a -.2 percentage point decrease in labor force participation, and a 2.3% decrease in log labor income. For the other outcomes, there are not statistically significant effects, but the point estimates go in the same direction as the results on labor income rank, labor force participation, and log labor income. Specifically, a one percentage point increase in physician opioid prescribing rates is associated with a -.065 (se=.048) percentile decrease in labor income rank for those with positive labor force participation, a .05 (se=.08) percentage point increase in probability of receiving any sick pay, a .15 (se=.08) percentage point increase in the probability of receiving more than two weeks of sick pay, and a .01 percentage point increase in probability receive DI.

These result suggests that the prescribing rate of opioids affects the extensive margin of labor supply, but it is unclear how important the effect on intensive margin is. The point estimate for the regression of labor income rank for those with positive labor force participation (.065) suggests that approximately half of the total effect on labor income rank (.12) is due to the intensive margin; however, the standard errors are large enough that we cannot rule out that there is no effect on the intensive margin of labor supply.

Panels B-D show the results for anti-inflammatories, anti-anxieties, and anti-depressants. We find no statistically significant effects on any of these other outcomes for any of the other drugs. For anti-anxieties, we continue to see positive effects on labor supply measures, but none of them are statistically significant at the 10% level. Importantly, for anti-depressants, even though there was a statistically negative effect on labor income rank (likely due to pre-period trends), we do not see any statistically significant effects on the other outcomes. The point estimate on labor force participation, labor income, and labor income rank conditional on positive participation, as well as DI point towards a negative effect on labor supply, though the effects on sick pay point toward a small positive effect.

2.5.3 Heterogeneity in Treatment Effects by Distance of Move

As a robustness check, we test for heterogeneity in the effect by distance of the move. Individuals who move longer distances are more likely to separate from their prior physician and thus have a larger change in prescribing rates. We therefore test whether they also have stronger effects on their drug use and labor supply, since this would provide evidence that the changes in labor supply and drug use are directly related to the changes in physician prescribing rates rather than differences in the effect of move for different individuals. Specifically, this relaxes the assumption that the correlation between the change in unobservables and individuals prior physician's prescribing rate is the same for individuals who move and do not move, and instead assumes that the correlation between the change in unobservables and individual's prior physician's prescribing rate is the same for individuals who move different distances. We find results that are consistent with our previous estimates, which suggests that the effects we see on labor income are due to the changes in physician prescribing rates.

Our distance measure is the change in probability other individuals who move between the same two municipalities see their pre-period physician after the move. This estimated separation rate we call: $\hat{s}_i$. To show that this distance measure affects the first stage, we bin the separation rate into 20 equal sized bins and estimate the coefficients in the following equation:

$$\hat{p}_{j(iA)} - \hat{p}_{j(iB)} = \alpha + \sum_{b=0}^{20} (\beta_b \hat{p}_{j(iB)} \mathbb{I}_{\hat{s}_i \in bin_b} + \pi_b \mathbb{I}_{\hat{s}_i \in bin_b}) + \epsilon_i \tag{2.11}$$

Where $\hat{p}_{j(iB)}$ is the individual's physician's prescription rate prior to moving, and $\mathbb{I}_{\hat{s} \in bin_b}$ is an indicator variable for whether the estimated separation rate from i's move is within bin_b. Figure 2.11a plots the coefficients $\hat{\beta}_b$ by the mean value of the separation rate in bin_b for prescribed opioids and shows a clear relationship between the magnitude of the change in prescription drug rates and the probability that individuals separate from their previous physician based on their origin and destination municipality.

To now understand how this affects an individual's probability they use the prescription drug, we estimate the following equation:

$$Drug_{it} = \alpha + \sum_{b=0}^{T} (\beta_b \hat{p}_{j(iB)} \mathbb{I}_{\hat{s}_i \in bin_b} + \pi_b \mathbb{I}_{\hat{s}_i \in bin_b}) + \sum_{b=0}^{T} (\mu_b \hat{p}_{j(iB)} \mathbb{I}_{\hat{s}_i \in bin_b} After_{it} + \rho_b \mathbb{I}_{\hat{s}_i \in bin_b} After_{it}) + \epsilon_i \tag{2.12}$$

Figure 2.11: *Heterogeneity in Effect of Physician Opioid Prescribing Rates by Distance of Move*

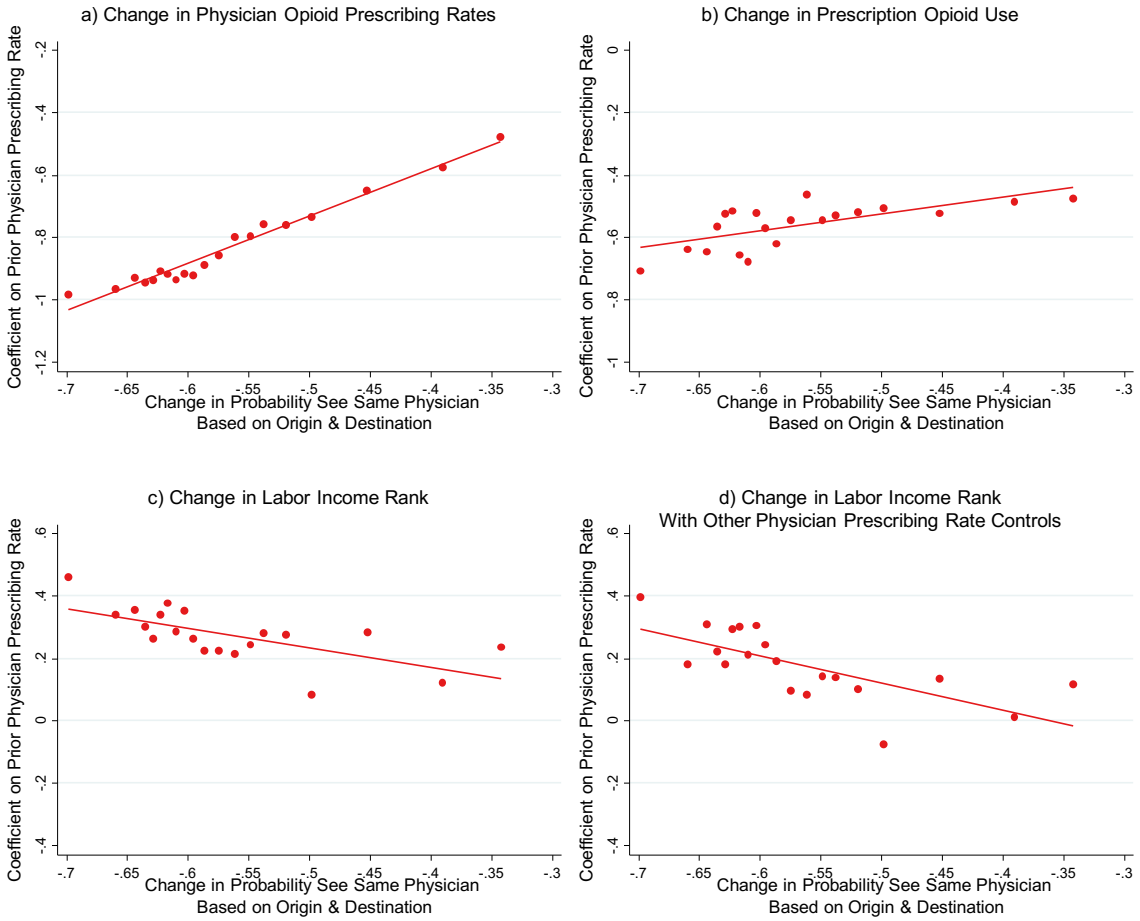

Notes: Panel (a) plots the coefficients of a regression of the change in physician prescribing rate on the pre-period physician's prescribing rate for different binned values of the change in the probability individuals see their same physician after their moved, which is based on the origin and destination of their move. Specifically, the change in the probability individuals see their same physician is calculated for all movers like in Figure 2.3(a), and we take the average difference in the post-period and the pre-period for all individuals who have the same origin-destination pair (or vice versa). We then bin these separation rates into 20 equal sized bins. We then estimate a separate coefficient on the pre-period physician's prescribing rates for each bin in a regression on the change in physician opioid prescribing rates- plotting the coefficients against the mean value of the separation rate. Note that this done only for the moving sample. Panel (b) uses the same x-axis, but we instead plot the coefficients of a regression of individual drug use on the pre-period physician's prescribing rate interacted with an after indicator and indicators for the different bins of the change in probability that see the same physician. Panel (c) is the same as Panel (b) except we use labor income rank as the outcome variable. Here we also include our standard set of controls for individual pre-characteristics (see Figure 2.7 Panel b) which we estimate separate coefficients on by bin. Panel (d) replicates (c) but includes controls for the other physician prescribing rates.

Since we previously found that individuals with certain characteristics had different effects of the move and the change in physician prescribing rates, we also allow for heterogeneity in the effects by age, gender, education, and previous blue collar status.

The coefficients of interest are μ_b,which are the coefficients on the interaction of the previous

physician's prescribing rate, an indicator for the separation rate being in bin b, and an indicator of after. They show how the relationship between the change in individual's prescription drug use and the pre-period physician's prescribing rate changes for different bins of the separation rate.

Figure 2.11b plots $\hat{\mu}_b$ by the mean value of $\hat{s}_i$ in bin_b for opioid prescription drug use. The figure shows that the relationship between the change in individual's drug use and the prior physician prescribing rate decreases for longer distance (ie higher separation rate) moves. Given that relationship between the change in physician prescribing rate and the previous physician's prescribing rate also decreases with longer distance moves (Figure 11a), this finding supports the evidence from Section 4.2, except here, instead of comparing the response between non-movers and movers, we compare the response between individuals who are more or less likely to be separated from their physician based on the distance of the move.

Now we move to the effect on labor income rank. We estimate the following equation:

$$LaborIncome_{it} = \alpha + \sum_{b=0}^{T}(\beta_b \hat{p}_{j(i)} \times \mathbb{I}_{\hat{s}_i \in bin_b} + \pi_b \mathbb{I}_{\hat{s}_i \in bin_b}) + \sum_{b=0}^{T}(\mu_b \hat{p}_{j(i)} \times \mathbb{I}_{\hat{s}_i \in bin_b} \times After_{it} + \rho_b \mathbb{I}_{\hat{s}_i \in bin_b} \times After_{it}) + \epsilon_i \quad (2.13)$$

Figure 2.11c plots $\hat{\mu}_b$ by the mean value of $\hat{s}_i$ in bin_b for opioid prescription drug use. The figure shows that the relationship between the change in individual's labor income after the move and the prior physician prescribing rate increases for longer distance moves. Since the longer distance moves led to a larger decrease in the physician prescribing rates, this suggests that a decrease in physician prescribing rates of opioids leads to an increase in labor income rank. In Figure 2.11d we additionally include the physician prescribing rates of the other drugs and find little effect on the results. This supports the evidence from Section 5.2.

In Appendix Figures B.6-B.8, we show a similar set of plots for anti-Inflammatories, anti-Anxieties, and anti-Depressants. For each drug we find that larger distance moves leads to larger changes in physician prescribing rates, and a larger change in individuals own prescription drug use for a given level of pre-period physician's prescribing rate. Appendix Figure B.6, Panels C and D show that longer distance moves have no effect on the relationship between prior physician inflammatory prescribing rates and the change in labor income rank, whether or not we include controls for the other prescribing rates (Panel D) or do not (Panel C). Appendix Figure B.7 Panel D

shows that once we include controls for the other prescribing rates, a longer distance move leads to a smaller relationship between the prior physician's anti-anxiety prescribing rate and labor income. This is inline with the results from Section 5.2 and Figure 2.9c. For anti-depressants, in Appendix Figure B.8 Panels C and D we see that longer distance moves leads to larger decreases in labor income, though the relation is fairly noisy once we include controls for the other prescribing rates.

2.6 Conclusion

In this paper, we show the effects of a change in physician prescribing rates of four important and widely used classes of drugs used to treat musculoskeletal and mental health disorders (opioids, anti-inflammatories, anti-anxieties, and anti-depressants) on a variety of outcomes. We find that a one percentage point increase in physician prescribing rates leads to an increase of approximately .45 percentage points in individual's own usage of the prescription drug. From this, we calculate that approximately 30% of the total variation in physician prescribing rates is due to causal physician effects rather than selection of patients. We find that generally individuals who are older, who are women, who are less educated, and work in blue collar occupation have larger effects of physician prescribing rates on their own prescription drug use, but there is some heterogeneity by the specific drug.

When we look at labor market outcomes, we find that an increase in physician's opioid prescribing rate leads to a decrease in individual's labor income and labor force participation: a 1 percentage point increase in physician's opioid prescribing rate leads to a decrease of .12 percentiles of labor income and .15 percentage point decrease in labor force participation. We do not find consistent effects of physician prescribing rates of the other drugs on labor supply outcomes.

So far we have been careful not to attribute the labor supply effects directly to the effects on the changes in prescription drug use. This is because while physician prescribing behaviors do strongly effect individual drug use, the prescribing rates may be correlated with many other physician traits that also help determine labor supply. However, we now discuss our results under the possibility that the effects we see *are* caused by prescription drug use, to better understand the

implications of this interpretation.

If we control for the other physician prescribing rates, remember that a 1pp increase in the physician prescribing rate of opioids led to a .6pp increase in opioid use and a .12 decrease in labor income rank and a .2 decrease in their labor force participation. If we make the interpretation that physician prescribing rates affect labor income only through their effect on prescription drug use, then our estimates that control for the prescribing rates of other drugs suggest that if an individual went from not prescription opioids to taking prescription opioids, we would expect their labor income rank to decrease by 20 percentiles and their labor force participation to decrease by 33 percentage points.[40] As a comparison, Boscarino et al. 2010 find that an estimated one fourth of out-patients on opioid chronic pain therapy develop opioid dependence. Granted that the increase in usage may be for a short period and not necessarily in long term chronic pain therapy, and since not all opioid dependence would necessarily lead to a drop out of the labor force, we might expect a smaller effect. Given that opioid prescribing rate may be correlated with other physician characteristics, we think of this estimate as an upper bound of the negative effect opioids have on labor supply.

For this interpretation, it is also important to note that we estimate a local treatment effect for individuals who are influenced by their physician's prescribing rate. This may be a particularly important caveat for interpreting the effects of anti-depressants. While we might expect anti-depressants to have a positive effect on labor income, as they clinically have been shown to alleviate depression - which can be debilitating,[41] we find some results which suggest that they have a negative effect on labor supply or no effect.[42] However, during the time period we study, anti-depressants grew tremendously such that by 2010, 10% of the working population took them. In the United States, where anti-depressant use has also increased dramatically, a recent study showed that nearly 2/3 of patients diagnosis with depression were given a false positive diagnoses, and the vast majority were given medication (Mojtabi 2013). This suggests that the prescribing

[40]Note that this requires a large amount of extrapolation: a one standard deviation increase in physician prescribing rates of opioids leads only to a 1.1 percentage point increase in individual's own prescription opioid use.

[41]Note that the one medical control trial that we know of that randomly assigned an anti-depressant or a placebo and looked at labor outcomes (N=43) found a negative effect on hours worked in the 6 weeks of follow-up, but it was not statistically significant (Agosti, Stewart, and Quitkin 1991).

[42]Note that due to large standard errors we cannot rule out some positive effect.

rate is at a level such that the marginal patient does not benefit from anti-depressant medication.

Similarly, opioids have also had a large increase in usage over this period. Our results are consistent with the interpretation that physicians are over-prescribing opioids such that the marginal patient has a negative effect on labor supply from the drug. While this result is more in line with the view that opioids have negative effects on labor outcomes due to their adverse effects and addictive properties, recent work by Kilby (2016) finds a decrease in opioid use has a negative effect on some measures of labor supply. Specifically, she uses variation from changes in opioid restrictions using differential timing across states in the onset of the Prescription Monitoring Program laws, and finds that these laws decrease the use of opioids, and increase the number of absent days at work for individuals with workers' compensation injuries and those on short term disability with pain related diagnosis codes. One important difference in our work is that Kilby's (2015) sample conditions on employed individuals. We find large effects on the extensive margin, so it is possible that the unconditional results may be different. An alternative explanation is that our results may be due to other differences in physicians who prescribe opioids. Due to the uncountably many physician characteristics, some of which are unobserved, we are unable to rule out this explanation.

In future work (with Fadlon, Nielsen, and Van-Parys), we plan to estimate directly the effects that different physicians have on labor supply. We will correlate these measures with various physician characteristics to find what are the most common characteristics that lead physicians to have positive impacts on their patients labor supply, which we can compare to the effects of prescribing rates.

Chapter 3

Mortality Place Effects: Evidence from Exposure to Municipalities in Denmark[1]

3.1 Introduction

To what extent is mortality shaped by the places individuals lived? While much research shows that places have large differences in mortality rates (and life expectancy), it is unclear to what extent this is due to selection of individuals to places versus causal effects of the place. Using Denmark Administrative population data of 21 years of residency, which we can link to subsequent 10 year mortality rates, we identify exposure effects of municipalities on mortality. We find that living in a municipality with a 1 SD higher exposure mortality effect for 60 years increases an individual's 10 year mortality rate by 3 percentage points (pp).

We identify exposure effects by analyzing individuals who moved across municipalities between 1984-2004. The first strategy we use looks at individuals who move only once during this period and compares the effect of the origin's mortality rate versus their destination's mortality rate by the year that they move. By comparing movers who move to and from similar places but just differ in the *timing* of the move, we are able net out any selection effect that is based on individuals ever choosing to live in particular municipalities. The remaining concern is that there is selection based on timing. To understand how this biases our results, we include

[1]Co-authored with Torben Nielsen

extensive pre-1984 controls based on the individual's income, education, labor force participation, hospitalizations, and chronic condition diagnoses. We find that when we include these controls our results attenuate by approximately 30%. With the full set of controls we find that an additional year at a place with a 1pp point higher 10 year mortality rate leads to a .03pp increase in the individual's own 10 year mortality. We use a secondary method that compares the effect of an additional year at a place using the full sample population and find very similar results.

We find that the effects are largest for the younger cohorts. Because different cohorts have very different base levels of 10 year mortality from 2005-2014, we switch our estimation from a linear model to a logit model. We find that the exposure effect to the average municipality mortality rate is approximately linear in the individual's cohort. This is consistent with a stock model of health, such that municipalities have smaller effects on total health for individuals who are older. We also test for heterogeneity in the effect by the individual's gender and their education prior to 1984, and find no evidence for heterogeneous effects by these characteristics.

Next, we estimate specific exposure effects and selection effects for each municipality by regressing 10 year mortality on a set of indicators for ever having lived in a municipality, (which picks up fixed selection effects), and variables for the number of years lived in each municipality (exposure effects) during the 1984-2004 period. We use standard shrinkage methods to estimate the variance of these two effects. When we only include cohort by gender fixed effects, we find that the standard deviation for the selection effect is .9 percentage points, while the standard deviation of exposure effects is .057 percentage points. When we include the full set of controls, the standard deviation for the "selection effects" attenuates to .5 pp and the standard deviation of the exposure effects attenuates to .051 percentage points. An F-test for the joint significance of the exposure effects has a p-value <.000. We find that these two effects are negatively correlated. Note while the size of the exposure effects may seem small, they add up over a lifetime such that living in a municipality with a 1 SD higher exposure mortality effect for 60 years increases an individual's 10 year mortality rate by 3 pp.

Finally, we explore the potential mechanisms through which exposure to municipalities may affect mortality. In particular, we look at how intermediate outcomes of individuals change when they move to a place with a higher exposure mortality rate. We find that their labor force participation, income, probability of taking prescription drugs and visits to general practitioners

decrease, but their Charlson score (an indicator for comorbidities) and hospital visits increase. This is suggestive evidence that a decrease in healthcare and income may be reasons for the differences in municipality exposure effects.

This paper relates to previous work that studies geographical differences in mortality rates (Diez Rouz and Mair 2010, Kulkarni et al. 2011). While they find large differences across areas, causal effects have been difficult to document due to the sparsity of long run geographical residence data linked to mortality data and the inherent selection problems. Some recent work also looks at the differences in life expectancy rates by certain variables (e.g. race and income) across areas (Chetty et al. 2016). Other recent literature has focused on estimating the causal effects of places on health care (e.g. Skinner 2012, and Finkelstein et al. 2016)

The methods we use are similar to those used in Chetty and Hendren (2016a) and Chetty and Hendren (2016b). They use childhood exposure to different places in United States on intergenerational mobility.[2] When we explore the mechanisms through which exposure to places may affect mortality, we use strategies similar to a growing literature which looks at how outcomes change when individual move to places with different characteristics (e.g. Finkelstein et al. 2016, Chetty et al. 2013, and Laird and Nielsen 2016).

The paper proceeds as follows: Section 2 details the data we use and important variable definition. Section 3 outlines a theoretical framework. Section 4 estimates average 10 year mortality rates by municipalities. Section 5 estimates the effect of exposure to places with a higher average mortality rate on individual's own mortality. Section 6 estimates specific exposure effects for each municipality, and then explores the mechanisms through which they affect mortality. Section 7 concludes.

3.2 Data

Our source of data is Danish Administrative records which links records of taxes, hospitalizations, physician reimbursements, deaths, prescription pick-ups, and employers.

[2]One difference to note is that they are able to use variation between siblings by the age of the move. Due to the fact that we are looking at moves that happen when the individuals are adults, we are unable to restrict ourselves to within family variation.

The most important variables we use are municipality residences from 1984-2004 for cohorts 1910-1960 and date of death.

Our analysis focuses on individuals who move between municipalities between 1984-2004, though we use non-movers as comparison. We classify a move as when an individual lives in two different municipalities in two consecutive years.[3] We restrict our analysis to individuals who were alive in 2004, who we observe their residence fully from 1984-2004, and who we have data on their labor income in 1980-1983. This excludes individuals who didn't live in Denmark at all within 1980-1983, or who lived elsewhere at anytime between 1984-2004.

3.2.1 Variable Definitions

Pre-1984 Characteristics: While our main analysis only requires residence and death information, we create a series of controls based on individual's characteristics prior to 1984, the period when we use to observe their residence.

Prior Labor Income Rank: We measure labor income as taxable labor earnings. We take the sum of labor income from 1980-1983, and then convert the sum into percentile ranks within an individual's year of birth, and their gender using the full sample of the Danish population. In the case of ties, we define the rank as the mean rank for the individuals in that group. For example, if 20% of women in the year 2000 who were born in 1950 have zero income, then they would all receive a rank of .10. We follow Chetty et al. (2014) who also use this measure because it creates a measure of relative income that is comparable across ages and gender, that is not overly influenced by the tails but still keeps their cardinal ordering; and that includes individuals with zero labor income.

Labor Force Participation: We calculate labor force participation as whether the individual has positive labor income within a year. We then take their average labor force participation from 1980-1983.

Months of Education: We calculate the highest months of education received by 1983. This

[3]In 2007, Denmark changed municipality boundaries by merging some municipalities. After the reform, the number of municipalities went from 293 to 99. We use the new municipality definitions and construct residence based on the new municipality borders prior to 2007. Due to anonymity concerns, we merge municipalities with fewer than 5000 residents to a larger municipality. All of these municipalities happen to be islands, so we merge them with the municipality that they have ferry access to.

measure of months of education is a function of the highest degree received. If the highest level of education was received prior to 1971, then the information is missing. When we use months of education as a control, we set the missing observations to zero and interact with an indicator that education is missing. 4% of individuals who ever moved between 1984-2004 have education missing.

Non-acute Hospitalizations: We calculate the number of years between 1977-1983 the individual had an non-acute non-pregnancy hospitalization.

Acute Hospitalizations: We calculate the number of years between 1977-1983 the individual had an acute non-pregnancy related hospitalization.

Modified Charlson Score: Using diagnoses from 1977-1983, we calculate whether the individual ever had a positive Charlson comorbidity index score.

Alternatives Outcomes: to understand the potential mechanisms through which places may affect mortality, we look at a couple of additional outcomes to see how they change around the time of a move to a place with a higher exposure effect in Section 6.3. The outcomes we look at our the following:

Hospitalizations: The number of non-pregnancy related hospitalizations within a year from 1977-2013.

General Practitioner Visits: The number of charges to a GP physician within a year from 1990-2013.

Labor Force Participation: Positive labor income from 1980-2011.

Labor Income: Yearly taxable labor earnings from 1980-2011. It is winsorized at 99th percentile for anonymity reasons.

Charlson Score: The sum of the charlson morbidity index based on diagnoses from a given year from 1977-2013. Winsorized at the 99th percentile.

Prescription Drug Use: An indicator for taking a prescription drug within a year, defined from 1995-2015.

Table 3.1: *Summary Statistics*

	Non-Movers		Move Once		Move > Once	
	Mean	SD	Mean	SD	Mean	SD
	(1)	(2)	(3)	(4)	(5)	(6)
Age in 2005	64	12	59	11	55	9
Female	0.538	0.499	0.513	0.500	0.455	0.498
Pre-1984 Characteristics						
Years of Education	10.2	4.2	11.0	3.8	11.2	3.5
# Acute Hosp Visits	0.223	0.536	0.274	0.604	0.340	0.704
# Non-Acute Hosp Visits	0.346	0.700	0.338	0.692	0.336	0.695
Positive Charlson Score	0.033	0.178	0.029	0.167	0.025	0.156
Labor Force Participation	0.852	0.317	0.880	0.272	0.889	0.243
Labor Income Rank	0.509	0.277	0.526	0.294	0.492	0.305
Mortality Place Effect	-0.037	0.014	-0.035	0.017	-0.035	0.017
Die 2005-2014	0.246	0.430	0.178	0.383	0.129	0.335
# Municipalities Lived	1	0	2	0	3.517	0.896
Number of Individuals	1610493		399443		186255	

Notes: This table reports the mean and standard deviation for various variables for three different samples. All samples are restricted to individual born from 1910-1960 and who lived in Denmark through 1980-2004. The first sample in Columns (1) and (2) are individuals who do not move between municipalities from 1984-2004. The second sample in Column (3) and (4) is of individuals who moved between municipalities only once. The final sample in Columns (5) and (6) are individuals who moved between municipalities more than once. Year of education is based of the highest degree the individual earned by 1983. The number of acute and non-acute hospital visits are the total from 1977-1983. They do not include hospitalizations due pregnancy. A positive charlson score is whether an individual had a positive charlson score anytime from 1977-1983. The charlson score is based off of diagnoses for comorbidities which are come from the hospitalization data. Labor force participation is the average labor force participation for the individual from 1980-1983. Labor force participation is an indicator for whether the individual had positive labor earnings. Labor income rank is calculated over the average labor income earnings for individuals from 1980-1983 by cohort and gender. The mortality place effect is calculated based on Equation (1) and is based off 10 year mortality rates and is calculated in reference to Copenhagen. Die 2005-2015 is whether the individual died between 2005-2014. Finally, # of municipalities lived refers to the number of unique municipalities individuals lived in from 1984-2004.

3.2.2 Summary Statistics

Table 3.1 details summary statistics for three different samples among the set of individuals from cohorts 1910-1960 and who were alive and observed in Denmark from 1980-2004. Column (1) and (2) give then mean and standard deviation of the sample of individuals who never moved between 1984-2004, which consists of 1,610,493 individuals. Columns (3) and (4) give the mean and standard deviation for the sample of individuals who moved only once during 1984-2004, which consists of 399,443 individuals. Finally columns (5)-(6) give the mean and standard deviation for the sample of individuals who moved more than once during 1984-2004, which consists of 186,255 individuals. We use all three samples during our analysis.

The average age for movers in 64 in the non-mover sample and is considerably younger for the mover samples (moved ever - 59 and moved more than once 55). Non-movers also tend to be female - 54% of the non-mover sample is female, while 52% of the moved once sample is female, and only 46% of the moved more than once sample is female.

In terms of characteristics observed prior to the moving observation period, non-movers have lower observed levels of education. Despite being younger, individuals who moved during this period were previously had more hospital visits for acute conditions, but slightly less visits for non-acute conditions. Those who move are less likely to have a positive charlson score (indicator for having a comorbidity) though it is unclear whether this is explained by them being a younger sample. 88% of the moving sample has a positive labor force participation prior to 1984. Those that move once have a higher than average labor income rank (.526) while those that more multiple times have lower labor income rank (.492) prior to 1984.

The non-movers have the highest mortality rate between 2005-2014 - at 24.6%, while the individuals who moved once have a mortality rate of 17.8%, and those who moved multiple times have a mortality rate of 12.9%.

3.3 Theoretical Framework

We want to estimate the causal effect places have on mortality. Each place has an average mortality rate: $\bar{a}_m$. This mortality rate is the sum of two things:

$$\bar{a}_m = s_m + c_m$$

A selection effect, s_m, a causal effect, c_m.

We assume that the causal effect of places, c_m, is a function of the place, the number of years the individual spent there and individual characteristics.

$$c_m = f(m, y_m, X_{it})$$

If $\mathbb{E}(s_m|y_m) = 0$, then we can identify one aspect of the causal effect, $\frac{\partial c_m}{\partial y_m}$ by finding $\frac{\partial \bar{a}_m}{\partial y_m}$. We call this effect the exposure effect.

Our identification allows for individuals sorting into different municipalities; however, the obvious concern is that $\mathbb{E}(s_m|y_m) \neq 0$, that is, individuals sort into places for certain amounts of time as well. To deal with this concern, we can take the $\mathbb{E}(a_i|y_m, X_{i,B})$ - control for individual's pre-characteristics. Given these are predetermined, they do affect the exposure effect, but they capture some of the selection effect. If the estimates change drastically with the sets of predetermined controls, this suggests that $\mathbb{E}(s_m|y_m) \neq 0$, and it remains unclear whether there is additional sorting on unobserved characteristics in addition to the observed characteristics. In the future work, we plan to bound the size of the effects under various assumption based on the amount of sorting (in terms of exposure) of observables to unobservables.

Most models of health and behavior that would allow places to have exposure effects require a stock element, such that even conditional on current behavior, past behavior still affects future mortality.

There are many potential mechanisms through which places might have these type of exposure effects. One potential mechanism is differences in peers across places who affect individual's health related behaviors. Another mechanism are differences in health care provision. While Denmark has a national health care system, with most services free, there still are geographical differences in the amount of health care (and the type) of healthcare provided (Laird and Nielsen 2016). Differences in preventative care may be particularly important in terms of exposure effects. Another mechanism could be affects on income. While health care is predominantly free in Denmark, individuals with higher income may be able to afford healthier food, access to gyms

and other activity, and lower stress over financial security. Additionally, places may have greater concentrations of occupations that have lower mental and physical stress on the body. In Section 6 of the paper, we explore a couple possible mechanisms.

3.4 Estimating Municipality 10 Year Mortality Rates

We estimate average municipality 10 year mortality rates using a linear model and by taking out cohort by gender fixed effects. Specifically we estimate the equation:

$$10YrDeath_i = \alpha_m + \gamma_{cf} + \epsilon_i \tag{3.1}$$

Where α_m are indicators for the municipality of residence in 2004, γ_{cf} are cohort by gender fixed effects, and $10YrDeath_i$ is whether the individual died between 2005-2014.[4] Copenhagen is the base category, so the effects of the other categories are in comparison to Copenhagen.

Figure 3.1a maps $\hat{\alpha}_m$ by municipality in Denmark. There are two things to note: the highest mortality rates are clustered in the Capital Region immediately surrounding Copenhagen, as well as the Zealand Region. The lowest rates of mortality are in Central Denmark Region. There appears to be a large amount of spatial correlation in the mortality rates.

Table 3.1 reports the variance of these effects Figure 3.1b plots a histogram of $\hat{\alpha}_m$ that is population weighted by residency from 1984-2004 of cohorts 1910-1960, who are observed for the full period.

3.5 Exposure Mortality Place Effects

Empirical Strategy

To estimate exposure mortality place effects, we first restrict our sample to individuals who move at least once between 1984-2004 and were born in 1960 or before.[5]

[4] We use a linear specification for ease of interpretation. Later we will use a logit specification in order to estimate heterogeneity in exposure effects.

[5] We make the cohort restriction so that all individuals are old enough so that the pre-period controls are informative.

Figure 3.1: *Distribution of Mortality Effects*

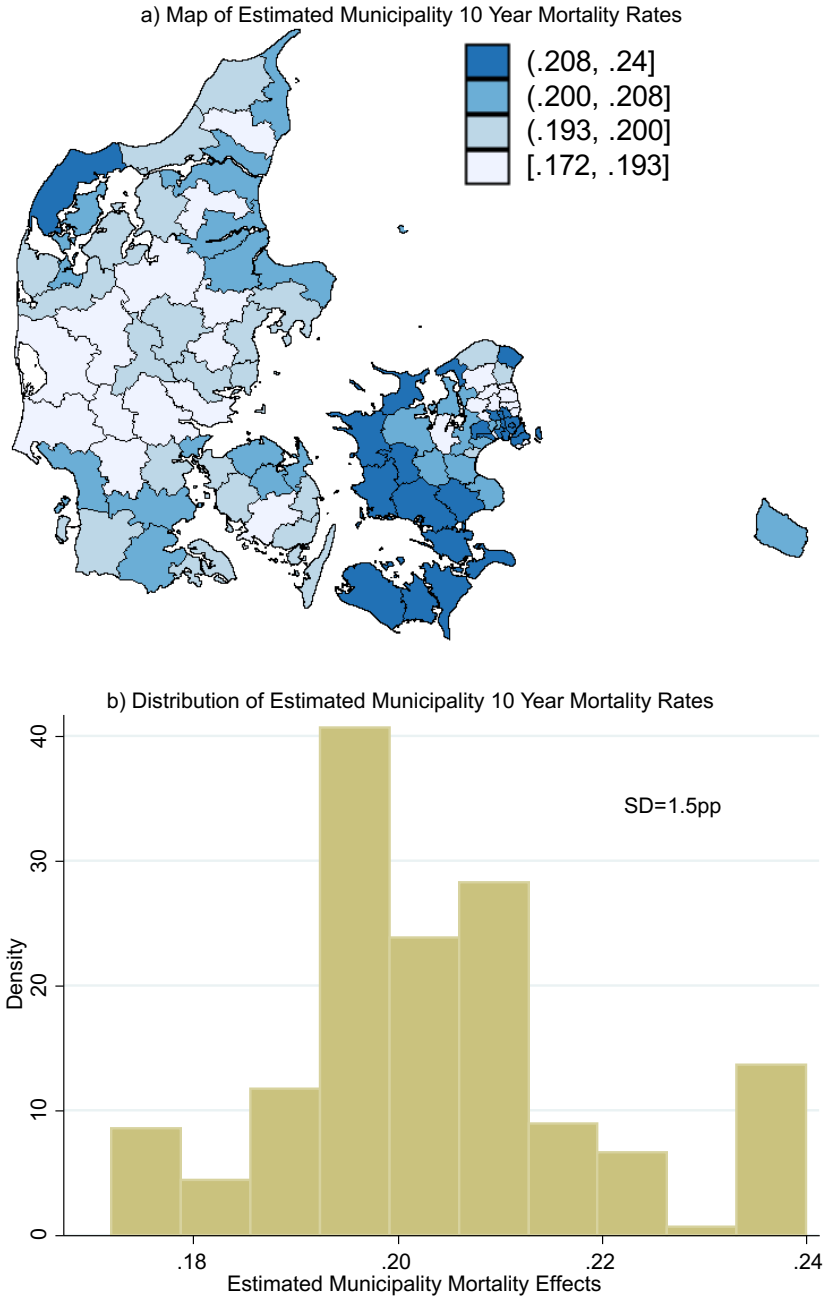

Notes: Figure 3.1a maps estimated municipality 10 year mortality effects, $\hat{\alpha}_m$. The mean 10 year mortality rate is added for reference. $\hat{\alpha}_m$ are estimate using a linear regression of dying between 2005-2014 on fixed effects for municipality residence in 2004 (α_m) and gender by cohort fixed effects. Figure 1b plots the histogram of these effects population weighted by residence in 2004.

The first method we use restricts the sample to individuals who moved once in that time period and estimate the effect of the destination's mortality rate versus the origin's mortality, by the year that the individual moved on whether the individual died from 2005-2015.[6]

To understand this empirical strategy, think of the following simple example: let us compare two individuals who are born in 1950. One moves to a high mortality place (H) from a low mortality place (L) at the end of the period, let say in 2004. On the other hand, the other individuals moves to H from L in 1985. The first individual has much more exposure to the high mortality place. If we there is a causal exposure effect, then we would expect that the first individual will be more likely to die in the next 10 years than the second individual.

To increase precision, the second method uses the full sample and estimates how the effect of a place's mortality rate changes based on the number of years the individual spent in that municipality between 1984-2004.

Identification Problems

There are a couple of identification problems with these methods that could bias our results.

Omitted Variable Bias

First, the second individual may be different from the first prior to the period, which led the individual to move to H later. To see if selection in terms of exposure time is driving the results, we include pre-period controls on the individual's income, labor force participation, education hospitalization record, and whether they have chronic illnesses to see how that changes the results. If the effects attenuate when we add these controls, that suggests that some of the effect we estimated is due to selection rather than a causal effect of the place.

Measurement Error

Second, we measure exposure based on residences from 1984-2004, which is a biased and noisy measure for total exposure prior to 2004. Our focus is on the comparison of an additional year spent in a place between 1984-2004, so it is important to understand how an additional year between 1984-2004 may reflect differences in exposure for the individual's entire lifetime. Unfortunately, we don't not observe the entire history residency for individuals; however, we

[6]We focus on the year that the individual moved rather than the age at which they moved since variation in the age is related to variation in the year of birth which affects the age at which we measure mortality.

can use the entire period we do observe residency (1980-2015) to understand how the last twenty years of residency reflects total residency over the entire time period. This is important because if later living choices are substitutes or complements to earlier in life choices, our estimates may over or underestimate the effect of an additional year in a place. To understand how exposure during 1984-2004 might reflect prior exposure, we using the full period of moves and estimate the effect of an additional year of exposure at a place with mortality rate $\hat{\alpha}_m$ between 1995-2015 on total experienced municipality based mortality rate ($\sum_{t=1980}^{2015} \hat{\alpha}_{m(i,t)}$ within the observable time period of 36 years (1980-2015).

Table 3.2: *Predictiveness of Residence from 1995-2015 on Total Residence 1980-2015*

	Yrs in Municipality 1980-2015		Total Muni Mortality Rate 1980-2015	
	(1)	(2)	(3)	(4)
Years in Municipality 1995-2015	1.038***	1.134***	0.035***	0.039***
	(0.001)	(0.001)	(0.0001)	(0.0001)
Years in Municipality 1995-2015 x Municipality 10 Yr Mortality Rate			0.930***	1.075***
			(0.004)	(0.003)
Number of Moves	1	>0	1	>0
N	744751	1365356	744751	1365356

Notes: The table estimates the effects of an additional year in municipality in 1995-2015 on the total number of years spent in the municipality from 1980-2015 in Columns (1) and (2). In columns (3) and (4) is estimates the effect of an additional year in a place with a particular mortality rate during 1995-2015 on the total mortality rate individual experienced (the average x 36). The odd columns are for individuals who moved between municipalities only once and the even columns are for individuals who ever moved between municipalities. The sample is restricted to individuals whose year of birth is between 1921-1971 and are present for the entire period.

In Table 3.2 columns (3) and (4), we report the coefficient π in the following regression:

$$\sum_{t=1980}^{2015} \hat{\alpha}_{m(i,t)} = \alpha + \beta NYrs9515 + \pi NYrs9515 \times \hat{\alpha}_m + \lambda \hat{\alpha}_m + \epsilon \tag{3.2}$$

When we restrict the sample to individuals who only moved once, we find that an additional year at a place with $\hat{\alpha}_m$ leads .93 increase in the total mortality rate experience by the individual,

suggesting there is some substitution. However, when we look at any movers, we find there is a small degree of complementarity such that an additional year at a place with $\hat{\alpha}_m$ leads 1.08 increase in the total mortality rate experience by the individual.

These results are suggestive that differences in exposure to places from 1984-2004 are associated with differences in total prior exposure that is close to, or perhaps slightly more than one on one.[7]

3.5.1 Results for Individuals who Move Once

To estimate the exposure of places on individuals, we first look at how the origin and destination municipality mortality rate effects an individual's own mortality based on the year at which an individual moves.

The specific equation we estimate is the following:

$$10YrDeath_i = \sum_{y=1984}^{2004} (\beta_y + \lambda_y^d \hat{\alpha}_d + \lambda_y^o \hat{\alpha}_o) + \tau X_{i1984} + \epsilon_i \tag{3.3}$$

Where X_{i1984} are various individual characteristics calculated from 1980-1983. In all specifications we include year of birth by gender fixed effects.

Figure 3.2 A plots the $\hat{\lambda}_y^d$ - the coefficients on the destination mortality rate in blue, $\hat{\lambda}_y^o$ - the coefficients on the origin mortality rate in red, by the year of the move. The figure shows that if the individual moved at the beginning of the period, then the coefficient on destination mortality rate is close to one while the coefficient on the origin mortality rate is close to zero. On the other hand, if the individual moved at the end of the period, the reverse is true. Since individuals who moved later spent more time in the origin relative to the destination, this suggests that places do have an exposure effect.

However, we still may be concerned that individuals who move in different years are different on other characteristics that lead to differences in their mortality. We therefore add a series of controls based on individuals characteristics measured prior to 1984. We plot the coefficients $\hat{\lambda}_y^d$ and $\hat{\lambda}_y^o$ when we include additional controls in Figure 3.2 Panel B. The effects attenuate by approximately 30%, but are still have significant slopes.

[7] If we restrict to individuals who only moved once from 1995-2015, we find that each additional year spent in a municipality from 1995-2015 is associated with 1.04 additional year spent in the municipality for the total period, which is closer to 1 on 1.

Figure 3.2: *The Effect of Destination vs Origin Mortality Rate on Own Mortality by Year of Move*

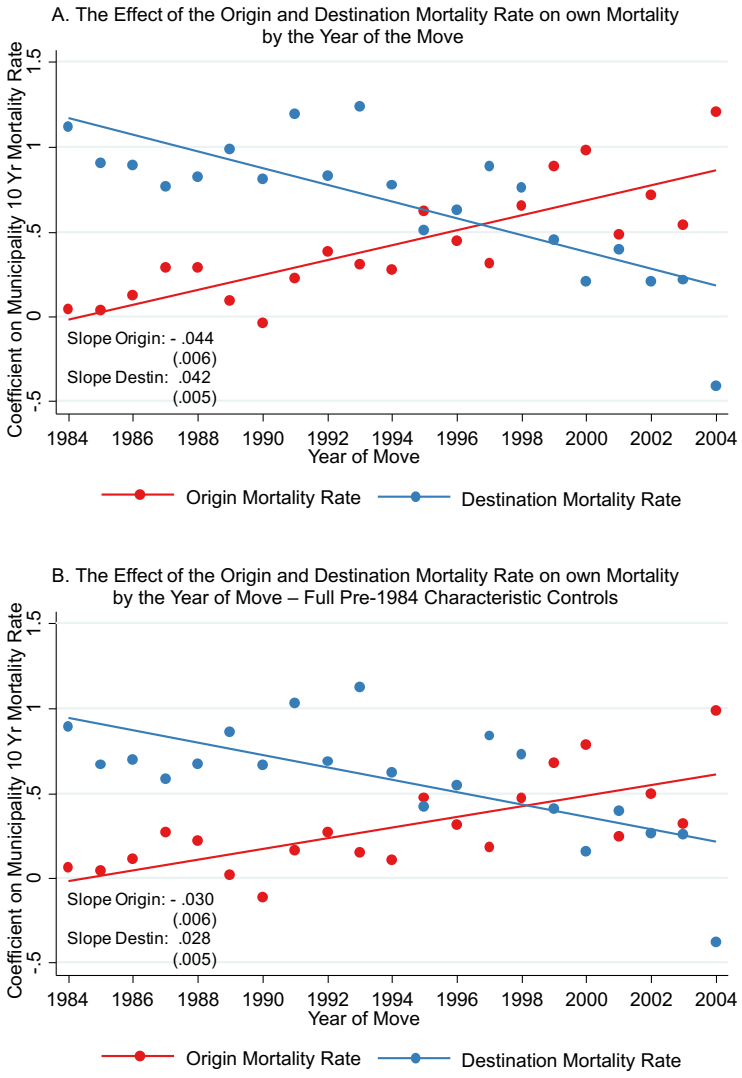

Notes: Figure 3.2a plots the coefficients from the regression (Equation 3.3) estimated off of individuals who move once between 1984 to 2004. Specifically, it plots the from the regression of individual 10 year mortality rate on indicators for the year of move interacted with the destination's estimated mortality effect ($\hat{\alpha}_d$) and indicators for the year of the move interacted with the origin's estimated mortality effect ($\hat{\alpha}_o$). For information on how $\hat{\alpha}_m$ are calculated see Figure 3.1 notes. Cohort by gender fixed effects are also included as controls for Panel A. For Panel B, we include additional controls based off of the individual's pre-1984 characteristics. Specifically we control for labor income rank, labor force participation, years of education, acute and non-acute hospitalizations, and charlson score indicator separately interacted with a quadratic in year of birth x gender. We report the slope coefficient based off the number of years of education from Table 3.3 Columns (1) for Panel A and (4) for Panel B.

In particular, in Table 3.3, we report the coefficients of the following regression which paramerizes years of the move linearly and includes various pre-1984 controls (X_{i1984}):

$$10YrDeath_i = \lambda^d(\hat{\alpha}_d \times YearMove_i) + \lambda^o(\hat{\alpha}_o \times YearMove_i) + \beta YearMove_i + \gamma^d\hat{\alpha}_d + \gamma^o\hat{\alpha}^o + \tau X_{i1984} + \epsilon_i \tag{3.4}$$

Column (1) includes just cohort by gender fixed effects like Figure 3.2A. Column (2) adds as controls the number of hospitalizations for acute conditions (from 1977-1983) interacted with quadratic in year of birth and gender, number of hospitalizations for non-acute conditions (from 1977-1983) interacted with quadratic in year of birth and gender, and average charlson score from 1977-1983 interacted with year of birth and gender. Column (3) adds as controls the average labor force participation 1980-1983 interacted with a quadratic in year of birth and gender, average labor income rank 1980-1983 interacted with a quadratic in year of birth and gender. And Column (4) add as controls the number of months of education interacted with a quadratic in year of birth and gender.

Column (1) shows that moving to a place with a 1pp point higher mortality rate one year later decreases one's own mortality by .04 percentage points, and alternatively moving *from* a place with a 1pp point higher mortality rate one year later increases one's own mortality by .04 percentage points. When we include additional controls for pre-period characteristics the effects attenuate. The largest change is when we include income and labor force participation controls in Column (3), with the coefficient of $\hat{\lambda}^d$ decreasing in magnitude from -.042 to -.034, and $\hat{\lambda}^o$ decreasing from .039 to .031.

Once we include the full set of of controls, we find that moving to a place with a 1pp point higher mortality rate one year later decreases one's own mortality by .03 percentage points, and moving *from* a place with a 1pp point higher mortality rate one year later increases one's own mortality by .028 percentage points. These estimates suggest that exposure to a place for 33 years leads to a one for one increase in one's own mortality rate based on the place's mortality rate.

While the effects attenuate when we include additional controls, there is still a robust exposure effect of places to individual's 10 year mortality. In particular, we find that living in a place with a 1 percentage point higher 10 year mortality rate for one additional year leads to a .03pp increase

Table 3.3: *The Effect of the Destination and Origin Muni. Mortality Rate by Year Moved on Mortality*

	10 Year Mortality			
	(1)	(2)	(3)	(4)
Destination Muni 10 Yr Mort Rate	1.098***	1.057***	0.934***	0.868***
	(0.0558)	(0.0557)	(0.0556)	(0.0556)
Destination Muni 10 Yr Mort Rate	-0.0436***	-0.0416***	-0.0335***	-0.0304***
x (Year Moved-1984)	(0.00560)	(0.00558)	(0.00556)	(0.00556)
Origin Muni 10 Yr Mort Rate	0.0130	0.0224	0.0268	0.0212
	(0.0480)	(0.0479)	(0.0478)	(0.0477)
Origin Muni 10 Yr Mort Rate	0.0415***	0.0392***	0.0316***	0.0284***
x (Year Moved-1984)	(0.00488)	(0.00487)	(0.00485)	(0.00485)
Cohort by Gender Fixed Effects	x	x	x	x
Pre-1984 Characteristics				
Hospital & Chronic Condition Controls		x	x	x
Income & LFP Controls			x	x
Education Controls				x
N	367651	367651	367651	367651

Notes: This table presents the coefficients estimated from Equation (3.4). It is restricted to the sample of individuals who moved between municipalities once between 1984-2004. The coefficients are from an OLS regression of an indicator for whether individual died between 2005-2014 on the destination municipality's 10 year mortality rate (taking out cohort by gender fixed effects), the interaction of the destination's 10 year mortality rate and the year moved, the origin municipality's 10 year mortality rate, the origin's 10 year mortality rate interacted with the year moved, and cohort by gender fixed effects and various other controls. Column (1) only includes cohort by gender fixed effects Columns (2)-(4) include additional controls based off of individual's pre-1984 characteristics (see Table 1 notes for definitions). Column (2) includes acute and non-acute hospitalizations and charlson score indicator separately interacted with a quadratic in year of birth x gender. Column (3) additional includes labor income rank and labor force participation separately interacted with year of birth x gender. Column (4) additionally includes years of education interacted with quaratic of year of birth x gender.

in an individual's own 10 year mortality rate. This confirms that the effect remains linear by year of move with controls.

3.5.2 Results for all Individuals

The previous method conditioned on moving only once in the 1984-2004 period, which severely restricts the sample. Here we use a secondary method that looks at the exposure of a place for *all* individuals.

Specifically, we run the following regression:

$$10YrDeath_i = \sum_{n=1}^{21} \gamma_n \hat{\alpha}_m + \tau X_{i1983} + \epsilon_{im} \tag{3.5}$$

Where γ_n estimates the effect of individual i being in municipality m for n years during the period 1984-2004. Note that γ_{21} is estimated off of individuals who don't move. We include the same set of controls as in Section 5.2.

Figure 3.3A plots γ_n when we only include cohort by gender fixed effects. The blue dot signifies the coefficient on the mortality rate for individuals who did not move during the period ($n = 21$). We also report the intercept (γ_0) and slope (γ) from the following regression, which are also reported in Table 3.4. In this regression we do not include the non-movers ($n = 21$).[8]

$$10YrDeath_i = \gamma \hat{\alpha}_m \times NYrs_{im} + \gamma_0 \hat{\alpha}_m + \beta NYrs_{im} + \tau X_{i1983} + \epsilon_{im} \tag{3.6}$$

These figures and tables show that an additional year spent in a place with a higher mortality effect of 1pp leads to a .033pp once we include the full set of controls. While there is some attenuation when we include the controls, the effect remains robust and large. Another thing to note is that the effect appears to be linear in the number of years spent in the municipality. This is additional evidence that places have causal exposure mortality effects on residents.

[8]This is partly because Figure 3.3 shows that the relationship between total number of years spent in the municipality from 1980-2015 and the number spent from 1995-2015 is discontinuous at year 21.

Figure 3.3: *The Effect of Place Mortality Rate by Number of Years Lived There*

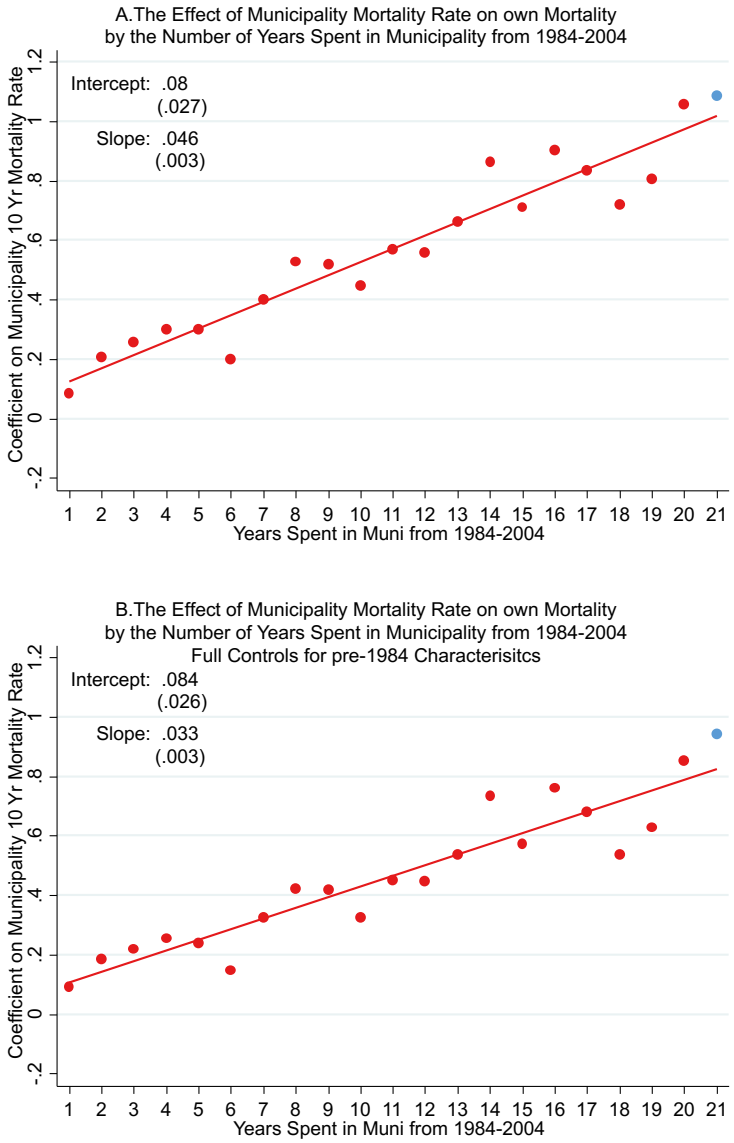

Notes: Figure 3.3a plots the coefficients from the regression (Equation 3.5) estimated off of individuals who we observe in Denmark from 1980-2004. Specifically, it plots the from the regression of individual 10 year mortality rate on indicators for the number of years lived in a municipality interacted with the estimated mortality effect ($\hat{\alpha}_m$) as well as cohort by gender fixed effects. The blue dot indicates the effect of the municipality mortality effect for individuals who lived in the municipality for 21 year - ie the entire period. For information on how $\hat{\alpha}_m$ are calculated see Figure 3.1 notes. For Panel B, we include additional controls based off of the individual's pre-1984 characteristics. Specifically we control for labor income rank, labor force participation, years of education, acute and non-acute hospitalizations, and charlson score indicator separately interacted with a quadratic in year of birth x gender. We report the slope coefficient based off the number of years of education from Table 3.3 Columns (1) for Panel A and (4) for Panel B.

Table 3.4: *The Effect of the Interaction of Muni. Mortality Rate and Municipality Exposure on Mortality*

	10 Year Mortality			
	(1)	(2)	(3)	(4)
Municipality 10 Yr Mortality Rate	0.0796**	0.0832**	0.0992***	0.0836**
	(0.0265)	(0.0264)	(0.0263)	(0.0263)
Municipality 10 Yr Mortality Rate x # of Years in Municipality	0.0462***	0.0436***	0.0364***	0.0332***
	(0.00258)	(0.00257)	(0.00256)	(0.00255)
Cohort by Gender Fixed Effects	x	x	x	x
Pre-1984 Characteristics				
Hospital & Chronic Condition Controls		x	x	x
Income & LFP Controls			x	x
Education Controls				x
N	1455201	1455201	1455201	1455201

Notes: This table presents the coefficients estimated from Equation (3.6). It is restricted to the sample of individuals who moved between municipalities at least once between 1984-2004. The unit of observation is individual municipality. The coefficients are from an OLS regression of an indicator for whether individual died between 2005-2014 on the municipality's 10 year mortality rate (taking out cohort by gender fixed effects), the interaction of the municipality's 10 year mortality rate and the number of years spent in that municipality between 2005-2014, and cohort by gender fixed effects. Column (1) only includes cohort by gender fixed effects Columns (2)-(4) include additional controls based off of individual's pre-1984 characteristics (see Table 1 notes for definitions). Column (2) includes acute and non-acute hospitalizations and charlson score indicator separately interacted with a quadratic in year of birth x gender. Column (3) additional includes labor income rank and labor force participation separately interacted with year of birth x gender. Column (4) additionally includes years of education interacted with quaratic of year of birth x gender.

3.5.3 Heterogeneity in Exposure Effects

We would like to know whether some individuals are more affected by place effects than others. This is difficult to answer with our current linear set-up since individuals may differ on their baseline mortality as well. We therefore switch our estimation equation to a logit model on 10 year mortality.[9]

We first reestimate α_m, using the logit model:

$$Pr(10YrDeath_i = 1|x_i) = \frac{exp(\alpha_m^L + \gamma_{cf})}{1 + exp(\alpha_m^L + \gamma_{cf})} \tag{3.7}$$

Where α_m^L are indicators for living in municipality m in 2004, γ_{cf} are cohort by gender fixed effects, and $10YrDeath_i$ is an indicator for individual i dying from 2005-2014.

After we estimate $\hat{\alpha}_m^L$, we then reestimate Equation 3.1 as a logit function and using $\hat{\alpha}_m^L$ instead of $\hat{\alpha}_m^L$:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\sum_{n=1}^{20} \gamma_n \hat{\alpha}_m^L + \tau X_{i1983} + \epsilon_{im}) \tag{3.8}$$

We do this to verify that when we use a logit model, we still find an exposure effect of estimated mortality place effects based on the number of years we observe their residence there from 1984-2005.[10] Figure 3.4a plots $\hat{\gamma}_n$ by the number of years lived at the place when we include the full set of pre-1984 controls, which include: income, labor force participation, education, hospitalizations controls and age by gender fixed effects. Again we see a strong linear increase in the effect of $\hat{\alpha}_m^L$ by the number of years an individual lived there. Table 3.5 column (1) reports the estimate of the effect of the γ when we measure number of years an individual lived somewhere continuously:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\sum_{n=1}^{20} \gamma \hat{\alpha}_m^L \times Nyrs_m + \nu Nyrs_m + \kappa \hat{\alpha}_m^L + \tau X_{i1983} + \epsilon_{im}) \tag{3.9}$$

[9] We use the linear model as the base model for ease of interpretation and because later it is easier to estimate the variance of effects.

[10] We restrict here to individuals who move.

Figure 3.4: *Heterogeneity of Effects and Logit Model*

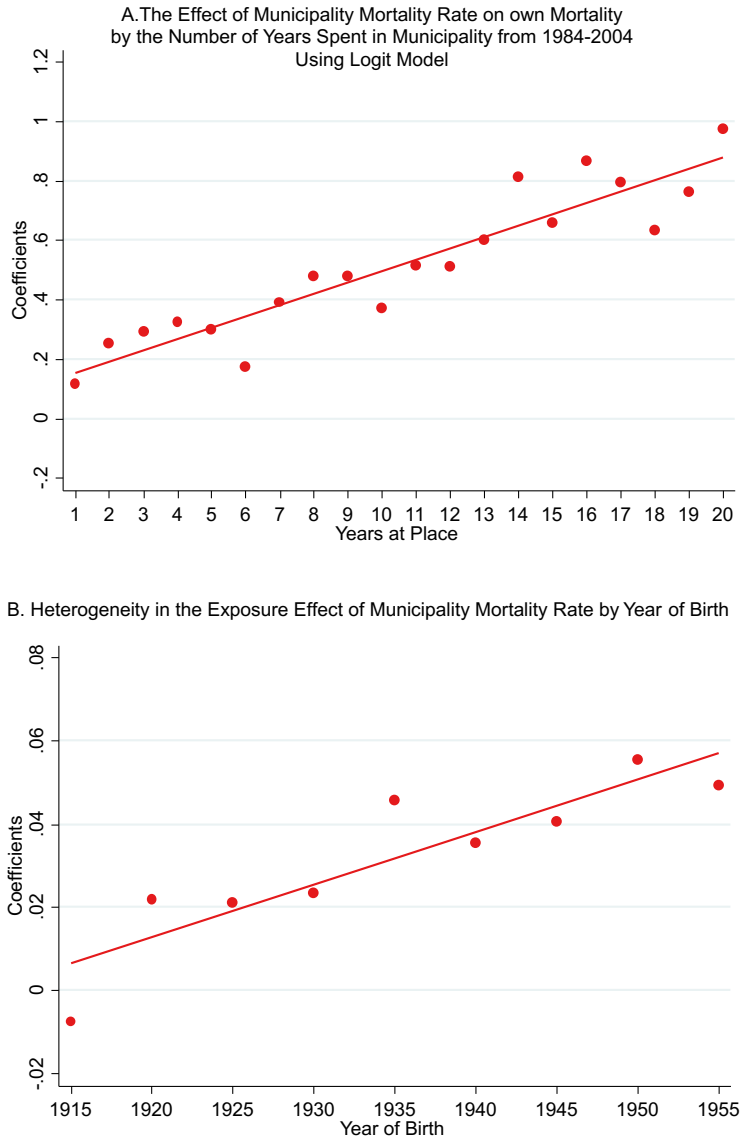

Notes: Figure 3.4a plots the coefficients from the logistic regression (Equation 3.8) estimated off of individuals who moved between municipalities in Denmark between 1984-2004. Specifically, it plots the from the logistic regression of individual 10 year mortality rate on indicators for the number of years lived in a municipality interacted with the estimated logistic mortality effect ($\hat{\alpha}_m^L$) as well as cohort by gender fixed effects and the full set of pre 1984 characteristic controls (see Figure 3.3B notes for details). $\hat{\alpha}_m^L$ is estimated from a logistic regression of 10 year mortality on municipality fixed effects and gender by cohort fixed effects. Panel 3.4b plots the coefficients from the logistic regression (Equation 3.9) of 10 year mortality on the full interaction of indicators for 5 year cohort bins with the number of years spent in a municipality and the municipality logistic mortality effect - $\hat{\alpha}_m^L$. It also includes the full set of pre-1984 controls.

Table 3.5: *Heterogeneity in the Exposure Effect*

	10 Year Mortality			
	(1)	(2)	(3)	(4)
Municipality Mortality Rate	0.113***	0.125**	-0.00815	0.349***
	(0.0320)	(0.0432)	(0.0974)	(0.104)
Muni Logit Mortality Effect	0.0382***	0.00581	0.0543***	0.0377***
x # of Years in Municipality	(0.003)	(0.00962)	(0.00932)	(0.00411)
Muni Logit Mort. Effect x # of Yrs in Muni		0.00103***		
x (Year of Birth-1910)		(0.000283)		
Muni Logit Mort. Effect x # of Yrs in Muni			-0.0112	
x Female			(0.00615)	
Muni Logit Mort. Effect x # of Yrs in Muni				0.00252
x Above Median Education				(0.00610)
Full Controls	x	x	x	x
N	1453744	1453744	1453864	1453864

Notes: This table estimates heterogenous effects of exposure effects by various year of birth, gender and above median education. Column (1) shows the estimates from Equation 3.8 which is a logit on municipality logit mortality effect and the interaction of the municipality logit mortality effect and the number of years the individual spent there between 1984-2004. It includes only movers. The municipality logit mortality effect is based of the coefficients from a logistic regression of 10 year mortality rate on municipality fixed effects and gender by cohort fixed effects. We include the full set of pre-1984 characteristics (those included in column (4) of Tables 3.3 and 3.4. In column (2), we add the full set of interaction between the municipality logit morality effect the number of years spent in the municipality, and the year of birth. For column (3), we instead include the full set of interaction between the municipality logit morality effect the number of years spent in the municipality, and an indicator for female. In Column (4), we instead include the full set of interaction between the municipality logit morality effect the number of years spent in the municipality, and having above the median level of education within one's cohort and gender group.

To see if there is heterogeneity in the effect by year of birth and the ages we measure residency, we first bin year of birth into approximately 5 year bins.[11] We then estimate the following:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\sum_{c=1915,int=5}^{1955} \pi_c \hat{\alpha}_m^L \times Nyrs_m + \nu Nyrs_m + \kappa \hat{\alpha}_m^L + \tau X_{i1983} + \epsilon_{im}) \quad (3.10)$$

Where π_c is the slope of $\hat{\alpha}_m^L$ for individuals in cohort bin c. Figure 3.4b plots π_c (there is no left out category) by the year of birth bins. We find that the exposure effects seem to increase with the year of birth, that is decrease with the age at which we measure individual's residency and subsequent mortality. We find close to no effects (with negative point estimate) for individuals born from 1910-1919, or for individuals who were 65-74 at the start of the observation of their residency. One interpretation is that individual's health is vastly predetermined by those ages, such that there are small to no exposure effects.

Table 3.5 Column (2) reports the coefficient, when we estimate continuous heterogenous effect in terms of year of birth:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\pi(Yob_i - 1910) \times \hat{\alpha}_m^L \times Nyrs_m + \chi Nyrs_m \times \hat{\alpha}_m^L + \kappa \hat{\alpha}_m^L + FOI + SOI + \tau X_{i1983} + \epsilon_{im}) \quad (3.11)$$

Where FOI and SOI refer to all other first order and second order interactions. We estimate that a one year increase in year of birth increases the size of the exposure effect by .001.

We also estimate whether there is differential effect by gender, by estimating the following equation:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\pi^f Female_i \times \hat{\alpha}_m^L \times Nyrs_m + \chi^f Nyrs_m \times \hat{\alpha}_m^L + \kappa^f \hat{\alpha}_m^L + FOI + SOI + \tau X_{i1983} + \epsilon_{im}) \quad (3.12)$$

The estimates for π^f, χ^f, and κ^f are reported in Table 3.5, column (3). We find that there is no

[11] We bin all individuals born from 1910-1919 since there are too few within the 1910-1914 bin. We also add individuals born in 1960 to the 1955 bin.

significant differential exposure for effect for women as opposed to men.

Finally, we estimate whether there is a differential effect by education. We create a variable $HighEd_i$ which is whether the individual has above the median level of education for their cohort and gender in 1983 prior to the observed residency period. We do this since some individuals may have not finished their education by 1983 and this differs by year of birth.

We estimate the following regression:

$$Pr(10YrDeath_i = 1|x_i) = Logit(\pi^e HighEd_i \times \hat{\alpha}_m^L \times Nyrs_m + \chi^e Nyrs_m \times \hat{\alpha}_m^L + \kappa^e \hat{\alpha}_m^L + FOI + SOI + \tau X_{i1983} + \epsilon_{im}) \quad (3.13)$$

The estimates for π^e, χ^e, and κ^e are reported in Table 3.5, column (3). We find that there is no significant differential exposure for effect for those who are highly educated and those who are less educated.

3.6 Municipality Specific Exposure Effects

3.6.1 Estimation of Municipality Specific Exposure Effects

We would like to know the specific exposure effect for each municipality. Due to the potential covariance of selection and exposure effects, we estimate them specifically rather than in relation to the average mortality rate.

For each individual who moves at least once between 1984-2004, we run a regression of:

$$10YrDeath_i = \sum_m (s_m EverLived_m + c_m \#Years_m) + \tau X_{i1983} + \epsilon_i \quad (3.14)$$

Where s_m are the coefficients on indicators for ever living in municipality m during 1984-2004, and c_m are the coefficients on the number of years lived in municipality m. We again condition on individuals whose residence we observe for the entirety of 1984-2004, and who we have pre-1984 controls for. We estimate equation 13 using different controls to see how our estimates change in order to understand how they maybe biased by time-varying selection.

3.6.2 Variance of Exposure Effects

To estimate the variance of the selection effect, captured by $var(s_m)$, and the variance in exposure effects, captured by $var(c_m)$, we use standard shrinkage methods to estimate the signal standard deviation of both effects since $\hat{s}_m$ and $\hat{c}_m$ are measured with error. Table 3.6 reports the signal

Table 3.6: *Signal Standard Deviation of Selection and Exposure Effects*

	Signal Standard Deviation				
	(1)	(2)	(3)	(4)	(5)
Regression 1					
Ever Lived in Municipality Effect	0.0174	0.0093	0.0082	0.0089	0.0079
Regression 2					
Ever Lived in Municipality Effect	0.0152	0.0085	0.0059	0.0078	0.0052
Exposure Municipality Effects of 1 Additional Year	0.00122	0.00057	0.00054	0.00055	0.00051
Cohort by Gender Fixed Effects		x	x	x	x
Income, LFP, and Educ Controls			x		x
Hospital & Chronic Condition Controls				x	x

Notes: This table reports the signal standard deviation of different estimated effects. The results under Regression 1 are from a regression of 10 year mortality on indicators for ever having lived in each municipality during 1984-2004. Regression 2 additional includes a set of variables for the number of years lived in each municipality. We estimate the signal standard deviation by subtracting the variance of the estimated noise of the estimates from the full variance of the estimated effects. Both regressions are estimated off the sample of individuals who moved at least once between 1984-2004. In Column (1) we include no additional controls. In Column (2) we include gender by cohort fixed effects. In columns (3)-(5) we additional control flexibly for pre-1984 characteristics. Specifically, in Column (3) we additionally control for labor income rank, labor force participation, and years of education separately interacted with year of birth x gender. Column (4) additionally includes acute and non-acute hospitalizations and charlson score indicator separately interacted with a quadratic in year of birth x gender. Column (5) includes both the social economic and health controls.

standard deviation of $\hat{s}_m$ and $\hat{c}_m$. The different columns report the standard deviations when we include various controls. Column (2) add year of birth by gender fixed effects. Column (3) additionally adds average labor income rank and labor force participation from 1980-1983 and months of education in 1983 interacted fully with a quadratic of year of birth and gender. Column (4) instead of including income, labor force participation, and education controls, it includes number of hospitalizations for acute and non-acute conditions from 1977-1983 and average Charlson score from 1977-1983 interacted with year of birth and gender. Column (5)

includes both the income controls from Column (3) and the hospitalization controls from column (4). If the exposure effect we are capturing is due to selection, we would expect that the variance in the effects to decrease when we include pre-period controls.

When we include no controls in Column (1) the standard deviation of selection effects is 1.52pp, and the standard deviation for exposure effects is .12pp. The selection effect measure the effect of ever living in a municipality and the exposure effect measures the effect of living an additional year in a place on 10 year mortality. We find that including cohort by gender fixed effects has a large effect on the standard deviation of both the selection (s_m) and the exposure effects (c_m) - decreasing the standard deviation of the selection effect to .85pp, and the exposure effect to .057pp. When we include controls for pre-1984 characteristics, the standard deviation of selection effects attenuates by 39% while the standard deviation exposure municipality effects attenuates by only 11%.

To interpret the standard deviation of our final estimate of the standard deviation of exposure municipality effects, we find that living in a municipality with 1 SD higher exposure effect for 10 years increases individuals 10 year mortality by .5 percentage points or 3% of the mean mortality rate for this sample. Note that a 60 year old individual who spent their entire life in a 1 SD higher mortality place would have a 10 year mortality rate that 3pp higher than a 60 year old who lived in an average place. For comparison, the standard deviation of municipality 10 year mortality rates (selection+exposure) is 1.7 percentage points.

3.6.3 Mechanisms

To understand the potential mechanisms through which the estimated exposure effects affect mortality, we study how individual's intermediate outcomes change when they move to a place with a higher exposure mortality rate.

We look at six main outcomes: individual labor force participation and income from (1980-2011), whether they took any drugs within a year (1995-2015), the number of charges to general practitioners (1990-2015), whether an individual has a positive Charlson score (an indicator for comorbidities 1977-2013), and the number of inpatient hospital visits that are unrelated to pregnancy (1977-2013).

Figure 3.5: *Map of 1 Year Exposure Effects on 10 Year Mortality*

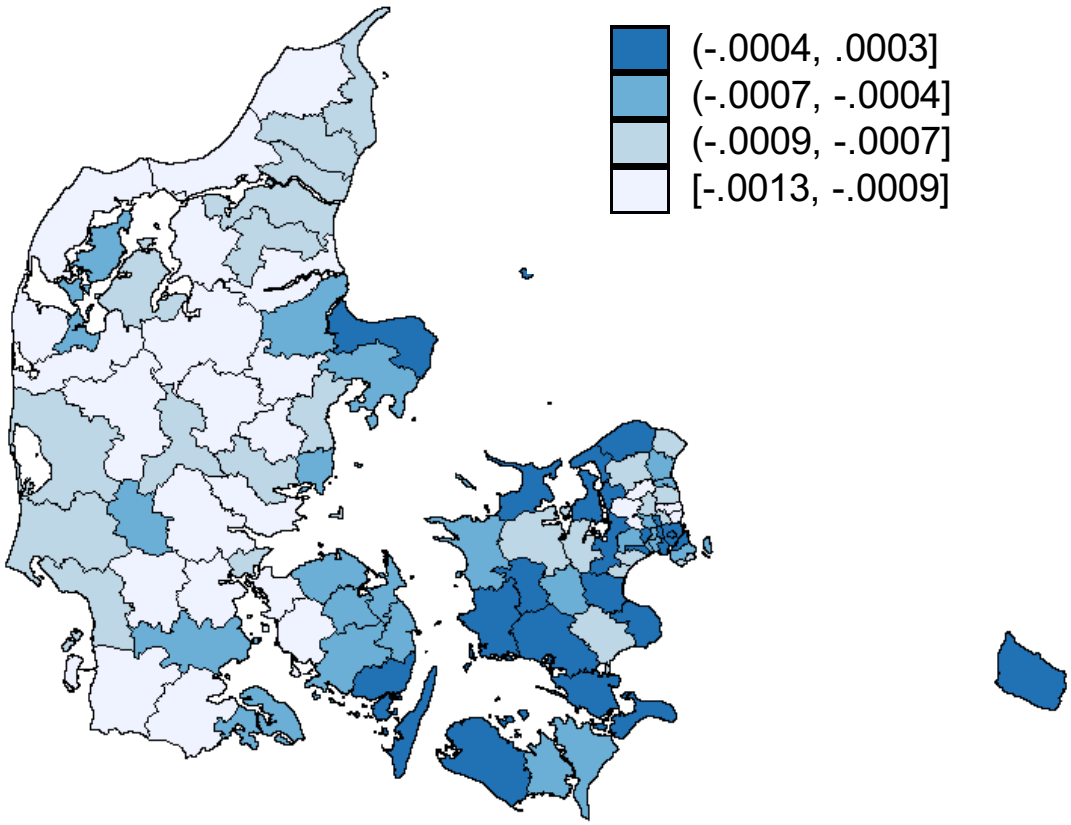

Notes: This figure maps the shrunk one year exposure effects of each municipality on ten year mortality. Specifically, they are the shrunk coefficients estimated from Equation 3.14 - that is the coefficients in the regression of 10 year mortality on the set of variables for the number of years spent in each municipality between 1984-2004. The regression also includes indicators for ever having lived in a municipality between 1984-2004, gender by cohort fixed effects and the full set of pre-1984 controls. See Figure 3.2 notes for details on the controls.

Specifically we estimate:

$$Y_{it} = \sum_{s=-3, s\neq -1}^{T} \theta_{r(i,t)}(\hat{c}^s_{d(i)} - \hat{c}^s_{o(i)}) + \pi(\hat{c}^s_{d(i)} - \hat{c}^s_{o(i)}) + \tau X_{it} + \mu_i + \epsilon_{it} \tag{3.15}$$

Where $\hat{c}^s_{d(i)} - \hat{c}^s_{o(i)}$ is the difference in estimated and shrunk exposure mortality effects between the destination and origin municipality that we estimated in Section 6.1, and $\theta_{r(i,t)}$ is a flexible function allowing for separate coefficients on $(\hat{c}^s_{d(i)} - \hat{c}^s_{o(i)})$ for each year relative to the move. We normalize θ_{-1} equal to zero so that the other coefficients indicate the effect of the change in exposure effects $(\hat{c}^s_{d(i)} - \hat{c}^s_{o(i)})$ relative to the year prior to the move. Finally, we also include age by gender, year, and year since event fixed effects (X_{it}).

We consider all moves during the observable period (1980-2015) in which the individual did not move in the observed pre or post period. We use 6 years of post data for all outcomes, except for prescription drug use, because prescription drug use data is restricted from 1995-2015. We do not balance the figures by individual, but we find similar results if we do.

Figure 3.6a shows that a move with a 1pp increase in the municipality exposure effect rate leads to a 3pp decrease in the labor force participation. After the third year after the move labor force participation recovers slightly so that by the end of 5 years, the individual's labor force participation is not significantly different from it was after the move.

Figure 3.6b looks at the effects of labor income, here reported in terms of 10K USD. We find that individuals who move to municipalities with a higher mortality rate have a sharp decrease in their labor income even *prior* to their move. This suggests that there is selection to higher mortality places for people who have a decrease in their income. After the move this trend continues, though the rate of decrease slows after the year of the move. Due to the pre-trend, it is difficult to know the causal effect that a place with a higher mortality rate has on individual's labor income. More investigation is required to understand what is going on.

Figure 3.6c shows that individual who move to place with a 1pp higher mortality exposure effect have a decrease in probability they take any prescription drug within a year of approximately 7 percentage points. A decomposition of this effect (not shown here), indicates this is largely due to decreases in anti-biotic and NSAID use (too of the most used drugs), though we do see significant decreases in other types of drugs (including opioids).

Figure 3.6: *The Effect of Moving to a Place with a Higher Mortality Exposure Effect on Intermediate Outcomes*

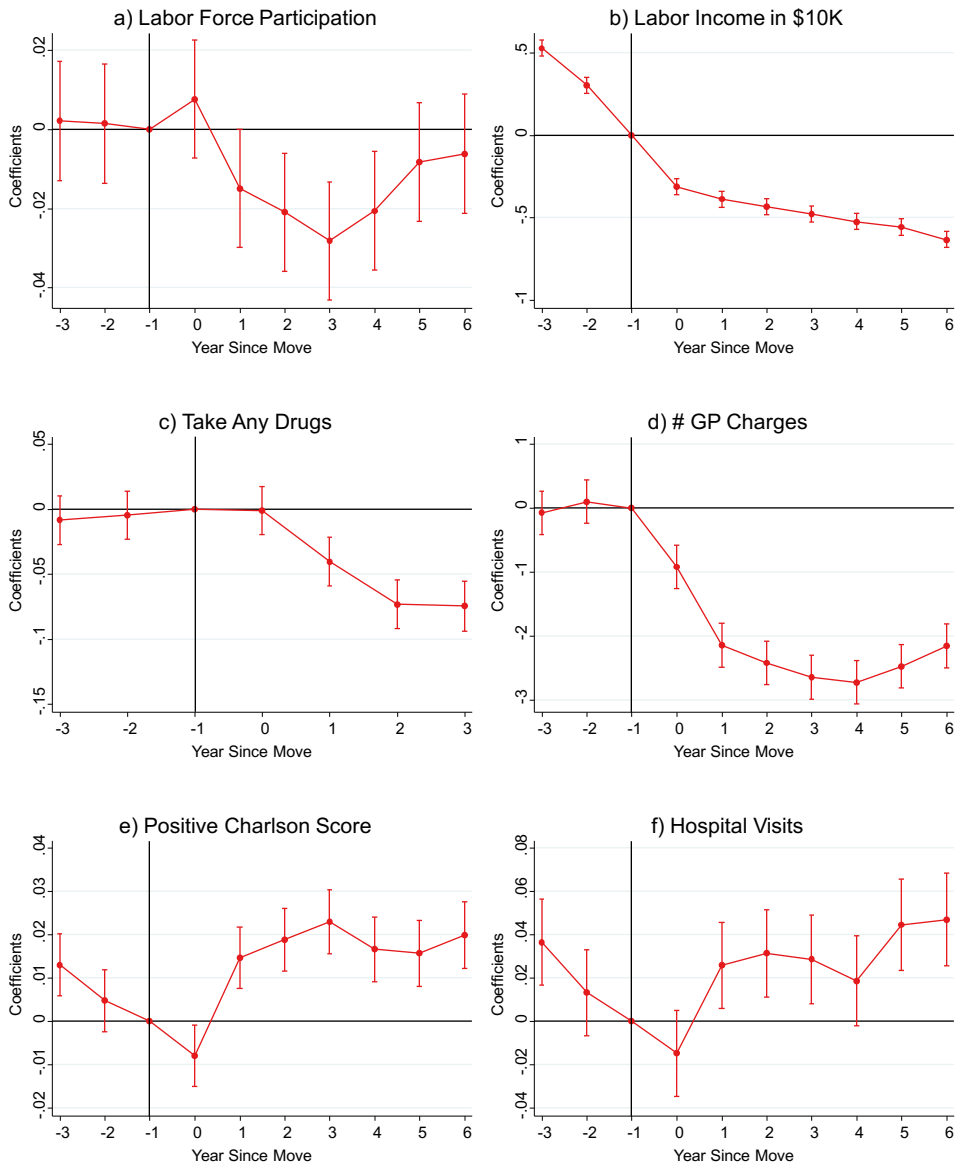

Notes: Figure 3.6 plots the effect of moving to a municipality with a shrunk mortality exposure effect that is 1pp higher. Specifically it plots the coefficients on the year since an individual moved interacted with the change in shrunk municipality mortality exposure effects. (The exposure effects are multiplied by 100 so that the coefficients can be interpreted as the effect of a 1pp point change in the exposure effect). See notes from Figure 3.5 for how the shrunk municipality exposure effects are estimated. The estimates are all in reference to the effect of the change in shrunk effects in the year prior to the move. 95% confidence intervals are plotted. Age by gender fixed effects and individual fixed effects are included in the regression. We consider potentially all moves during the observable period (1980-2015) in which the individual did not move in the observed pre or post period. We use 6 years of post data for all outcomes, except for prescription drug use (panel c).In Panel A the outcome is labor force participation defined from 1980-2011. In Panel B the outcome is labor income in $10,000 defined from 1980-2011. In Panel C the outcome is whether the individual took any prescription drug within the year which is defined from 1995-2015. In Panel D the outcome is the number of charges to a General Practitioner within a year. In Panel E the outcome is a positive charlson score which is an indicator for a co-morbidity. This is based off of diagnoses from hospitalization data, which is defined from 1977-2013. In Panel F the outcome is the number of inpatient hospital visits within a year that are unrelated to pregnancy, which is defined from 1977-2013.

Figure 3.6d shows that individuals who move to a place with a 1pp higher mortality exposure effect have a decrease in the number of their general practitioner charges by approximately 2.5.

Figure 3.6e shows that individuals who move to a place with a 1pp higher mortality exposure effect have an increase in the probability of having a positive charlson score (an indicator for comorbidities) by approximately two percentage points. There is a small *decreasing* pre-trend that breaks only in the year after the move. We find very similar results for the number of hospitalization visits (non-pregnancy related) - there is a decreasing trend prior to and through the move, and then an increase in hospitalization visits after the move for individuals who move to places with higher mortality exposure effects.[12]

The evidence from these event studies suggest that places with higher exposure mortality effects have causal negative effects on some types of health care received, and labor force participation. They also appear to have positive causal effects on receiving a comorbidity diagnosis and hospitalizations. While it is difficult to parse these effects as those due changes in health versus changes in availability of services, these results are consistent with possibility that labor market opportunities and health care (in form of drugs and physician visits) are important determinants of places effects on mortality. More research is required.

3.7 Conclusion

In conclusion, we find substantial exposure mortality effects of places. Individuals who live in places with higher mortality rates for a longer amount of time have higher mortality rates themselves. These effects are the largest for younger cohorts which is consistent with places being more important when individuals are younger and developing their health. When we estimate exposure effects specifically, we find that moving to a place with a 1SD higher mortality exposure effect for 10 years leads to an increase of ten year mortality of .5pp.

To understand the mechanisms involved, we look at the effect of moving to a place with a higher mortality exposure effect on intermediate outcomes. We find that individuals visit the GP less and take fewer prescription drugs after the move, but are hospitalized more. This is consistent

[12]Note that the similarity between the results for hospital visits and the charlson score is partly due to the fact that charlson score is based on diagnoses during hospital visits.

with higher mortality places having fewer preventative services, which leads to more extreme health outcomes later.

References

Aaronson, S. and Coronado, J. L. (2005). Are firms or workers behind the shift away from db pension plans? *FEDS paper*, (2005-17).

Agosti, V., Stewart, J. W. and Quitkin, F. M. (1991). Life satisfaction and psychosocial functioning in chronic depression: effect of acute treatment with antidepressants. *Journal of affective disorders*, **23** (1), 35–41.

Ashton, H. (2005). The diagnosis and management of benzodiazepine dependence. *Current opinion in Psychiatry*, **18** (3), 249–255.

Bengtsson, S. (2002). Bestemmer forvaltningen om du får førtidspension. *København: Socialforskningsinstituttet*, **2**, 15.

Berndt, E. R., Bailit, H. L., Keller, M. B., Verner, J. C. and Finkelstein, S. N. (2000). Health care use and at-work productivity among employees with mental disorders. *Health Affairs*, **19** (4), 244–256.

Bernheim, B. D. (1998). Financial illiteracy, education, and retirement saving. In O. Mitchell and S. Schieber (eds.), *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*, Philadelphia: Univ. Pennsylvania Press, pp. 38–68.

— and Rangel, A. (2009). Beyond revealed preference: Choice-theoretic foundations for behavioral welfare economics. *Quarterly Journal of Economics*, **124** (1), 51–104.

Beshears, J., Choi, J., Laibson, D., Madrian, B. and Weller, B. (2010). Public policy and saving for retirement: The "autosave" features of the pension protection act of 2006. In J. Siegried (ed.), *Better Living Through Economics*, pp. 274–290.

—, Choi, J. J., Laibson, D. and Madrian, B. C. (2009). The importance of default options for retirement saving outcomes: Evidence from the united states. In *Social security policy in a changing environment*, University of Chicago Press, pp. 167–195.

Bingley, P., Gupta, N. D. and Pedersen, P. J. (2007). Fiscal implications of reforms in retirement systems in denmark. In *Social Security Programs and Retirement around the World: Fiscal Implications of Reform*, University of Chicago Press, pp. 119–154.

Bjordal, J. M., Ljunggren, A. E., Klovning, A. and Slørdal, L. (2004). Non-steroidal anti-inflammatory drugs, including cyclo-oxygenase-2 inhibitors, in osteoarthritic knee pain: meta-analysis of randomised placebo controlled trials. *bmj*, **329** (7478), 1317.

Bollini, P., Pampallona, S., Tibaldi, G., Kupelnick, B. and Munizza, C. (1999). Effectiveness of antidepressants. meta-analysis of dose-effect relationships in randomised clinical trials. *The British Journal of Psychiatry*, **174** (4), 297–303.

Brobeck, S. (1990). *U.S. Consumer Knowledge: Results of a Nationwide Test*, , Washington, DC: Consum. Fed. Am.

Bütikofer, A. and Skira, M. (2015). Missing work is a pain: the effect of cox-2 inhibitors on sickness absence and disability pension receipt.

Cai, L. and Cong, C. (2009). Effects of health and chronic diseases on labour force participation of older working-age australians. *Australian Economic Papers*, **48** (2), 166–182.

—, Mavromaras, K. and Oguzoglu, U. (2014). The effects of health status and health shocks on hours worked. *Health economics*, **23** (5), 516–528.

Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C. and Metrick, A. (2009). Optimal defaults and active decisions. *Quarterly Journal of Economics*, **124** (4), 1639–1674.

Case, A. and Deaton, A. (2015). Rising morbidity and mortality in midlife among white non-hispanic americans in the 21st century. *Proceedings of the National Academy of Sciences*, **112** (49), 15078–15083.

Ceulemans, D., Hoppenbrouwers, M.-L., Gelders, Y. and Reyntjens, A. (1985). The influence of ritanserin, a serotonin antagonist, in anxiety disorders: a double-blind placebo-controlled study versus lorazepam. *Pharmacopsychiatry*, **18** (05), 303–305.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T. H. and Olsen, T. (2014a). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *Quarterly Journal of Economics*, **129** (3), 1141–1219.

—, — and Saez, E. (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *The American Economic Review*, **103** (7), 2683–2721.

— and Hendren, N. (2016). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects.

—, —, Kline, P., Saez, E. and Turner, N. (2014b). Is the united states still a land of opportunity? recent trends in intergenerational mobility. *The American Economic Review*, **104** (5), 141–147.

—, Looney, A. and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, **99** (4), 1145.

—, Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A. and Cutler, D. (2016). The association between income and life expectancy in the united states, 2001-2014. *Jama*, **315** (16), 1750–1766.

Choi, J. J. (2015). Contributions to defined contribution pension plans. *Annual Review of Financial Economics*, **7** (1).

—, Laibson, D. and Madrian, B. C. (2007). *Mental Accounting in Portfolio Choice: Evidence from A Flypaper Effect*. Tech. rep., National Bureau of Economic Research.

—, —, — and Metrick, A. (2003). Optimal defaults. *American Economic Review*, pp. 180–185.

—, —, — and — (2004). For better or for worse: Default effects and 401 (k) savings behavior. In *Perspectives on the Economics of Aging*, University of Chicago Press, pp. 81–126.

Chou, R., Turner, J. A., Devine, E. B., Hansen, R. N., Sullivan, S. D., Blazina, I., Dana, T., Bougatsos, C. and Deyo, R. A. (2015). The effectiveness and risks of long-term opioid therapy for chronic pain: a systematic review for a national institutes of health pathways to prevention workshop. *Annals of internal medicine*, **162** (4), 276–286.

Cremer, H., De Donder, P., Maldonado, D. and Pestieau, P. (2008). Designing a linear pension scheme with forced savings and wage heterogeneity. *International Tax and Public Finance*, **15** (5), 547–562.

— and Pestieau, P. (2011). Myopia, redistribution and pensions. *European Economic Review*, **55** (2), 165–175.

Currie, J. and Madrian, B. C. (1999). Health, health insurance and the labor market. *Handbook of labor economics*, **3**, 3309–3416.

—, Stabile, M. and Jones, L. (2014). Do stimulant medications improve educational and behavioral outcomes for children with adhd? *Journal of health economics*, **37**, 58–69.

Cutler, D., Stern, A. D., Skinner, J. and Wennberg, D. (2015). Physician beliefs and patient preferences: A new look at regional variation in health care spending.

Danish Ministries of Business and Economics, Finance, Employment and Taxation (2003). *Større valgfrihed I pensionsopsparingen*. Commisioned report.

Davis, P., Gribben, B., Scott, A. and Lay-Yee, R. (2000). The "supply hypothesis" and medical practice variation in primary care: testing economic and clinical models of inter-practitioner variation. *Social science & medicine*, **50** (3), 407–418.

Daysal, N. M. and Orsini, C. (2014). The miracle drugs: Hormone replacement therapy and labor market behavior of middle-aged women.

Dell'Osso, B. and Lader, M. (2013). Do benzodiazepines still deserve a major role in the treatment of psychiatric disorders? a critical reappraisal. *European Psychiatry*, **28** (1), 7–20.

Diez Roux, A. V. and Mair, C. (2010). Neighborhoods and health. *Annals of the New York Academy of Sciences*, **1186** (1), 125–145.

Dobkin, C., Finkelstein, A., Kluender, R. and Notowidigdo, M. J. (2016). *The Economic Consequences of Hospital Admissions*. Tech. rep., National Bureau of Economic Research.

Doyle, J. J., Ewer, S. M. and Wagner, T. H. (2010). Returns to physician human capital: Evidence from patients randomized to physician teams. *Journal of health economics*, **29** (6), 866–882.

Dummann, K. (2008). What determines supply and demand for occupational pensions in germany? *Journal of Pension Economics and Finance*, **7** (02), 131–156.

Fadlon, I. and Laibson, D. (2016). *Endogenous Paternalism and Pseudo-Rationality*. Unpublished manuscript.

— and Nielsen, T. H. (2016). Household labor supply responses to severe health shocks and the gains from social insurance. In *Social Insurance Programs (Trans-Atlantic Public Economic Seminar-TAPES)*, American Economic Journal: Economic Policy.

Finkelstein, A., Gentzkow, M. and Williams, H. (2016). Sources of geographic variation in health care: Evidence from patient migration. *The Quarterly Journal of Economics*, **131** (4), 168101726.

Garthwaite, C. L. *et al.* (2012). The economic benefits of pharmaceutical innovations: The case of cox-2 inhibitors. *American Economic Journal: Applied Economics*, **4** (3), 116–37.

Gaskin, D. J. and Richard, P. (2012). The economic costs of pain in the united states. *The Journal of Pain*, **13** (8), 715–724.

Gelber, A. (2011). How do 401(k)s affect saving? evidence from changes in 401(k) eligibility. *American Economic Journal: Economic Policy*, **3** (4), 103–122.

Gibbons, R. and Katz, L. (1992). Does unmeasured ability explain inter-industry wage differentials? *The Review of Economic Studies*, **59** (3), 515–535.

Goda, G. S. and Manchester, C. F. (2013). Incorporating employee heterogeneity into default rules for retirement plan selection. *Journal of Human Resources*, **48** (1), 198–235.

Grytten, J. and Sørensen, R. (2003). Practice variation and physician-specific effects. *Journal of health economics*, **22** (3), 403–418.

Gustman, A. L., Steinmeier, T. L. and Tabatabai, N. (2009). Do workers know about their pension plan type? comparing workers' and employers' pension information. *Overcoming the Savings Slump: How to Increase the Effectiveness of Financial Education and Saving Programs*, pp. 47–81.

Hastings, J. S., Madrian, B. C. and Skimmyhorn, W. L. (2013). Financial literacy, financial education, and economic outcomes. *Annual Review of Economics*, **5**, 347–73.

Hernæs, E., Piggott, J., Zhang, T. and Strøm, S. (2011). Occupational pensions, tenure, and taxes. *Journal of Pension Economics and Finance*, **10** (03), 435–456.

Hilgert, M. A., Hogarth, J. M. and Beverly, S. G. (2003). Household financial management: The connection between knowledge and behavior. *Fed. Res. Bull.*, **89**, 309.

Jäger, S. (2015). How substitutable are workers? evidence from worker deaths.

Kantor, E. D., Rehm, C. D., Haas, J. S., Chan, A. T. and Giovannucci, E. L. (2015). Trends in prescription drug use among adults in the united states from 1999-2012. *JAMA*, **314** (17), 1818–1830.

Kilby, A. (2016). Opioids for the masses: welfare tradeoffs in the regulation of narcotic pain medications. In *2016 Fall Conference: The Role of Research in Making Government More Effective*, Appam.

Kocsis, J. H., Zisook, S., Davidson, J., Shelton, R., Yonkers, K., Hellerstein, D. J., Rosenbaum, J. and Halbreich, U. (1997). Double-blind comparison of sertraline, imipramine, and placebo in the treatment of dysthymia: psychosocial outcomes. *American Journal of Psychiatry*, **154** (3), 390–395.

Krueger, A. (2016). Where have all the workers gone?

Kulkarni, S. C., Levin-Rector, A., Ezzati, M. and Murray, C. J. (2011). Falling behind: life expectancy in us counties from 2000 to 2007 in an international context. *Population health metrics*, **9** (1), 16.

Laird, J. and Nielsen, T. (2017a). Danish geographical variation in prescription drug use. *Working Paper*.

— and — (2017b). The effects of physician prescribing behaviors on prescription drug use and labor supply: Evidence from movers in denmark. *Working Paper*.

Li, S. and Laxminarayan, R. (2015). Are physicians' prescribing decisions sensitive to drug prices? evidence from a free-antibiotics program. *Health economics*, **24** (2), 158–174.

Lin, J., Zhang, W., Jones, A. and Doherty, M. (2004). Efficacy of topical non-steroidal anti-inflammatory drugs in the treatment of osteoarthritis: meta-analysis of randomised controlled trials. *Bmj*, **329** (7461), 324.

Liu, Y.-M., Yang, Y.-H. K. and Ou, H.-T. (2010). Hospital competition and prescribing behaviors: empirical evidence from taiwan. *The Journal of Developing Areas*, **43** (2), 125–141.

Lusardi, A. and Mitchell, O. S. (2007). Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, **54** (1), 205–224.

Lydiard, R. B., Stahl, S. M. and Harrison, W. M. (1997). A double-blind, placebo-controlled study comparing the effects of sertraline versus amitriptyline in the treatment of major depression. *The Journal of clinical psychiatry*, **58** (11), 484–491.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, **116** (4), 1149–1187.

Maestas, N., Mullen, K. J. and Strand, A. (2013). Does disability insurance receipt discourage work? using examiner assignment to estimate causal effects of ssdi receipt. *The American Economic Review*, **103** (5), 1797–1829.

Mercuri, M., Natarajan, M. K., Norman, G. and Gafni, A. (2012). An even smaller area variation: Differing practice patterns among interventional cardiologists within a single high volume tertiary cardiac centre. *Health Policy*, **104** (2), 179–185.

Mitchell, O. S. (1988). Worker knowledge of pension provisions. *Journal of Labor Economics*, **6** (1), 21–39.

Molitor, D. (2014). The evolution of physician practice styles: evidence from cardiologist migration. *University of Illinois at Urbana-Champaign Working Paper: http://www. business. illinois. edu/dmolitor/movers. pdf*.

OECD (2009). Pension country profile: Denmark. In *OECD Private Pensions Outlook 2008*, Paris: OECD Publishing.

O'Neill, L. and Kuder, J. (2005). Explaining variation in physician practice patterns and their propensities to recommend services. *Medical Care Research and Review*, **62** (3), 339–357.

Roeder, K. (2014). Optimal taxes and pensions with myopic agents. *Social Choice and Welfare*, **42** (3), 597–618.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, pp. 34–55.

Shaheed, C. A., Maher, C. G., Williams, K. A., Day, R. and McLachlan, A. J. (2016). Efficacy, tolerability, and dose-dependent effects of opioid analgesics for low back pain: A systematic review and meta-analysis. *JAMA internal medicine*.

Skinner, J. *et al.* (2012). Causes and consequences of regional variations in health care. *Handbook of health economics*, **2**, 45–93.

Stewart, J. W., McGRATH, P. J., Quitkin, F. M., Rabkin, J. G., Harrison, W., Wager, S., Nunes, E., Ocepek-Welikson, K. and Tricamo, E. (1993). Chronic depression: response to placebo, imipramine, and phenelzine. *Journal of clinical psychopharmacology*, **13** (6), 391–396.

—, Quitkin, F. M., McGrath, P. J., Rabkin, J. G., Markowitz, J. S., Tricamo, E. and Klein, D. F. (1988). Social functioning in chronic depression: effect of 6 weeks of antidepressant treatment. *Psychiatry research*, **25** (2), 213–222.

Thirumurthy, H., Zivin, J. G. and Goldstein, M. (2008). The economic impact of aids treatment labor supply in western kenya. *Journal of human resources*, **43** (3), 511–552.

Timbie, J. W., Horvitz-Lennon, M., Frank, R. G. and Normand, S.-L. T. (2006). A meta-analysis of labor supply effects of interventions for major depressive disorder. *Psychiatric services*.

Van Tulder, M. W., Scholten, R. J., Koes, B. W. and Deyo, R. A. (2000). Nonsteroidal anti-inflammatory drugs for low back pain: a systematic review within the framework of the cochrane collaboration back review group. *Spine*, **25** (19), 2501–2513.

Wang, Y. R. and Pauly, M. V. (2005). Spillover effects of restrictive drug formularies on physician prescribing behavior: Evidence from medicaid. *Journal of Economics & Management Strategy*, **14** (3), 755–773.

Appendix A

Appendix to Chapter 1

A.1 Verification of Employer Responses

Recall that contribution rates within a workplace, defined by all employees with the same 2-digit occupation code in the same firm, can vary since employers may set pension contributions at a finer level within the firm (i.e., higher-digit occupation codes), and since individuals have some ability to choose a different distribution of contributions between capital and annuity accounts than the default set by the employer. Our main analysis identifies employer default contribution rates using workplace-level medians.

To validate our measure of choice and to verify that our results are not driven by employees' responses within their employer-sponsored accounts, but rather by the employers themselves, we augment the analysis in the following ways:

1. **Analysis of Modes:** Instead of workplace-level medians we calculate workplace-level modes for identifying employer contribution rates. Panel A.1 of Appendix Figure 2 replicates Panel B of Figure 1, but with changes in workplace-level modes of contribution rates instead of medians, and shows that our results persist with this choice of the outcome variable. To further show that the patterns are very similar to those provided by medians, Panel A.2 of Appendix Figure 2 replicates Panel B of Figure 1 that uses medians, but for the same sample that is included in Panel A.1 of Appendix Figure 2. This sample includes only workplaces with unique mode values. Due to rounding issues and multiplicity of possible modes

we are left with 59,956 workplace-year observations, compared to 84,764 workplace-year observations in the main analysis of medians, which is the reason for our choice of medians over modes.

2. **Deviation from Aggregates:** Panels B.1 and B.2 of Appendix Figure 2 plot the distribution of the distance between employee-level capital contribution rates and workplace-level aggregates, for modes and medians, respectively. The analysis verifies that most employee contributions bunch exactly at these aggregates, supporting our design and choice of measures for identifying employer behavior.

3. **Workplace Size:** Statistically, medians are more likely to accurately identify default contribution rates in large workplaces. Conceptually, the firm decision making process is more likely to be tailoring defaults to groups of employees of similar occupations in larger workplaces. Therefore, we study the sensitivity of our analysis to the size of workplaces that are included in our sample. Appendix Figure 3 plots changes in medians for the samples of workplaces with more than ten, twenty, and fifty employees. While the sample size substantially decreases, the patterns of the results are very similar.

4. **Workplaces with No Contributions to Annuity Accounts:** To focus on a sample of employees with less discretion over contributions to employer-sponsored accounts, we repeat the analysis for employees who could not have reallocated contributions across different types of accounts. In particular, Panel C of Appendix Figure 2 constrains the sample to workplaces whose median annuity contribution rate in the years prior to the reform (1996-1998) was zero. In addition, we include only workplaces whose median capital contribution rate in 1998 was economically meaningful (here we choose having rates larger than 1.5%), to focus on employers that had non-negligible potential to reduce capital contributions in response to the reform. These restrictions substantially reduce the sample size (to a total of 4,056 observations), but reveal the same patterns.

A.2 Distribution of Employee-Level Changes in Contribution Rates

In section 1.3.2 we compared employer responses to the reform to the responses of individuals within their private accounts by using workplace-level aggregates. To account for differential baseline contribution rates to private vs. employer-sponsored accounts, we analyze percent changes in the two types of accounts at the employee level, by plotting the distribution of employee-level year-to-year percent changes in capital contribution rates. In addition, to conduct a within-employee analysis when we compare responses within private vs. employer-sponsored accounts, we balance the sample such that we keep only employees that had positive contributions to both type of accounts in a previous year.

In Appendix Figure 4 we plot histograms of individual-level percent changes in contributions to capital accounts for employer-sponsored capital accounts (Panel A) and private capital accounts (Panel B). In each panel we plot the distributions for the two years prior to the reform (so that we average the changes from 1996 to 1997 and from 1997 to 1998) and the distribution of changes following the reform – which captures changes from 1998 to 1999.

Comparing the before and after distributions, we see that in both types of accounts most of the effect of the reform is driven by exiting capital accounts altogether (the large differences in the mass points at -100 in both panels). In fact, in both employer-sponsored and private accounts, there was an increase of 23 pp in the fraction of employees who exited capital accounts.[1]

Recall from Section 1.3.2 that employer responses were much larger than individual responses in the aggregate. The evidence here, which results from analyzing the contributions of the special sample of individuals with positive lagged contributions to both accounts, suggests that at the individual level percent change responses in private accounts were similar to those in employer accounts. However, this result is not entirely surprising since the analysis here diverges from the ideal experiment that we mentioned is Section 1.3.2 for two main reasons. First, we analyze only a small share of individuals. While among those that were at the top tax bracket in 1999 55% had positive contributions to employer-sponsored accounts and 30% had positive

[1]Note, however, that while the responses to the reform within private accounts were composed of opting out of capital accounts by almost exclusively those who would otherwise have small increases or no changes in contributions, opting out of employer-sponsored accounts was also equally driven by workplaces who would have otherwise decreased their contributions (the decrease in Panel A from 10 pp to 3 pp in the mass point at -4% that includes changes larger or equal to -4% and smaller than 0%).

contributions to private accounts, less than 13% had positive contributions to both. Specifically, these workers are likely to be the most attentive to their pension accounts than a randomly selected individual who does not typically contribute to private pension accounts. Second, the sample includes affected individuals in workplaces that also had to cater to employees who were not affected by the reform. Therefore, changes in employer-sponsored accounts do not capture the full potential extent of employer responses. The average fraction of workers at the top income tax bracket in these workplaces was 74% in 1999, and given the linear response we found in Section 1.3.1, there is a potential attenuation of the employer responses by approximately 26% compared to the response of employers with all workers at the top income tax bracket.

A.3 Capital Contributions vs. Labor Income Responses

Since we are analyzing employer capital contributions as a fraction of labor income, the heterogeneous responses of firms by the fraction of workers above the top threshold could be due to differential changes in the numerator, employer capital contributions, or the denominator, taxable labor income. Appendix Figure 5 breaks down the responses in capital contribution rates plotted in Panel A of Figure 1 into annual earnings and capital contribution levels.

Using different workplace-level measures for annual labor earnings, Panel A, B, and C show no noticeable differential patterns over the years across workplaces with different shares of employees above the top tax threshold. In contrast, Panel D, which plots workplace-level medians of contribution levels, exhibits exactly the same patterns of Panel A of Figure 1, confirming that our results are driven by changes in employer contributions to pension accounts.

This analysis also allows us to test an implication of the standard model of employee compensation. Recall that the decrease in subsidies for capital contributions potentially led to an income effect, as it reduced the overall compensation of workers above the top tax threshold. However, Appendix Figure 5 as well as the time series of the log of labor income in Panel B of Appendix Figure 1, display no detectable change in compensation through increased annual earnings, either in the aggregate or as a function of the share of employees above the top tax threshold.[2]

[2]This is consistent with the required compensation of the loss incurred at the initial bundle of contributions (on

A.4 Inclusion of Workers Covered by Collective Agreements

Recall that our analysis aims at isolating workplaces in which default contribution rates are set by the employer. To this end, we excluded workers in the public sector or with blue-collar occupations (as they are likely covered by collective agreements), so that we included only workers in private firms with white-collar occupations. However, some white-collar jobs in the private sector are covered by collective bargaining.

In this appendix, we assess how inadvertently including workers covered by collective agreements may have affected our results. Inclusion of such workers can alter the results in two different ways. First, given that there are union representatives within the collective bargaining process, employer contributions for these workers may be more closely related to workers' preferences than employer contributions that are set exclusively by firms. Thus, employer capital contributions for them may respond more to the reform. In this case, inadvertently including workplaces that are covered by collective bargaining could increase the magnitude of the coefficients. On the other hand, in such workplaces, employer contributions are set for groups of occupation-level worker unions and employer associations, not at the finer occupation-firm level. Since we analyze the occupation-firm cell as the decision cell in our analysis, variation in the share of employees above the top tax threshold across employers but within the same worker union-employer association unit will not exhibit differential responses. This will tend to flatten the gradient of responses with respect to the share of affected employees and to attenuate our results.

In order to empirically quantify how unintentionally including workplaces covered by collective bargaining affects our estimates, we analyze the relationship between employer capital contributions and the fraction of workers above the top tax threshold for different groups that are either more or less likely to be covered by collective bargaining agreements as compared to workers with white-collar occupations in the private sector.

Specifically, in column (3) of Appendix Table 1 we restrict the analysis to employees in the

the order of no more than a few thousands of DKr) being negligible relative to annual labor earnings (on the order of hundred thousands of DKr). Empirically, such small changes are hard to detect, and, conceptually, some degree of wage rigidity due to, e.g., re-negotiation costs, would render these small changes non-profitable.

public sector or those with blue-collar occupations, who are more likely to be covered by collective bargaining. The coefficient of interest – on the interaction of the fraction of workers above the top threshold with the indicator for year 1999 – is -0.35, which is an order of magnitude smaller than (and statistically different from) the corresponding coefficient of -2.18 in our main sample analysis of employees with white-collar occupations in the private sector (replicated in column (1) of Appendix Table 1). In column (2) of Appendix Table 1 we restrict the analysis to private-sector workplaces of white-collar occupations with a highly-educated workforce. In particular, we focus on workplaces in which more than twenty percent of workers had more than sixteen years of education, which are even less likely to be covered by collective bargaining (as compared to all private-sector workplaces of white-collar occupations). The point estimate of the coefficient of interest in this case is large, -3.68 pp, and is significantly different from the coefficient for the remainder sample of private-sector white-collar occupations.

Put together, the results suggest that inadvertently including some workers who are covered by collective bargaining in our main specification may have attenuated our results.

A.5 Occupation Codes

In this appendix, we list 2-digit occupation codes used by Statistics Denmark. In our analysis, white-collar occupations are defined as occupations with codes whose first digit is between 1 and 5.

Occupation Codes:

1 Leadership at the top level of companies, organizations and the public sector

11 Legislative work and leadership in public administration and interest groups

12 Senior management of the company

2 Work that requires knowledge at the highest level in the area

21 Work within the non-biological branches of science and computer science, statistics, architecture and engineering sciences

22 Work in medicine, pharmacy, and biological branches of science, as well as midwives, general nursing work etc.

23 Teaching in primary schools, vocational schools, colleges, universities, and research organizations

24 Work within the social sciences and humanities

3 Work that requires Intermediate knowledge

31 Technicians in non-biological topics

32 Technicians and other work in biological topics

33 Caring and educational work

34 Work in sales, finance, business, administration, etc.

4 Office jobs

41 Internal office work

42 Office work with customer service

5 Retail sales, service and care work

51 Service and care work

52 Retail and models

6 Work in agriculture, horticulture, forestry, hunting and fishing

61 Work in agriculture, horticulture, forestry, hunting and fishing

7 Craft and related trades workers

71 Work in mining and construction

72 Metal and engineering work

73 Precision craftsmanship, graphic work, etc.

74 Other craft and work in related trades

8 Machine operating and work in transportation and civil engineering

81 Stationary plant

82 Operation of industrial machinery

83 Transportation

9 Other work

91 Cleaning and renovation work, messenger and security services, telephone canvassing

92 Assisting in agriculture, horticulture, fisheries and forestry

93 Manual work in the construction sector, manufacturing and transportation

A.6 Figures and Tables

Table A.1: *Changes in Employer Contribution Rates to Capital Pension Accounts for Different Samples*

Dependent Variable:	Δ Capital Contributions		
Sample:	White-Collar in Private Sector		
	All (1)	Highly-Educated (2)	Blue-Collar or in Public Sector (3)
Fraction of Employees Above Top Tax Threshold (Baseline Year 1998)	0.132	0.303	0.138**
	(0.145)	(0.611)	(0.074)
Fraction of Employees Above Top Tax Threshold Interacted With:			
Year 1996	0.273	-0.159	-0.093
	(0.167)	(0.752)	(0.081)
Year 1997	0.096	-0.234	0.063
	(0.147)	(0.675)	(0.076)
Year 1999	-2.182***	-3.681***	-0.347***
	(0.203)	(0.861)	(0.084)
Year 2000	-0.593***	-1.467**	-0.129*
	(0.173)	(0.747)	(0.077)
Year 2001	-0.553***	-0.178	-0.08
	(0.176)	(0.915)	(0.078)
Year Fixed Effects	X	X	X
Income and Workforce Size Controls	X	X	X
2-Digit Occupation-Firm Fixed Effects	X	X	X
2-Digit Occupation-Year Fixed Effects	X	X	X
Number of Observations	60,643	7,989	78,533
Number of Clusters	20,642	3,742	25,932

Notes: This table reports estimates of employers? responses to the reform, for different samples, as a function of the share of their employees whose earnings were above the top labor-income tax threshold. The outcome variable is the change in employer capital contribution rates from the previous year. We regress this outcome on the fraction of workers above the top tax threshold, year fixed effects, the fraction of workers above the top tax threshold interacted with year fixed effects, and a set of controls as indicated in the table. The baseline year is 1998, so that the coefficient on the fraction of employees above the top tax threshold refers to that year. The coefficient on the fraction of employees above the top tax threshold interacted with other year indicators estimates this relationship relative to the relationship in the baseline year. Income controls include a fifth-order polynomial of the mean workplace-level labor income, separately for workers above and below the top tax threshold, as well as their interactions with year indicators. Workforce size controls include the number of workers in a workplace and its square, as well as their interactions with year indicators. Column (1) replicates exactly column (4) of Table 2, which restricts the sample to private-sector firms and white-collar occupations. Column (2) repeats the analysis of column (1) for private-sector firms and white-collar occupations, but restricts the sample to workplaces with a highly-educated workforce. Specifically, we focus on workplaces in which more than twenty percent of workers had more than sixteen years of education. Column (3) restricts the estimation to the sample of employees in the public sector or those with blue-collar occupations. All the regressions exclude self-employed individuals and workplaces with less than five employees. The observation units are workplaces, defined as all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. Standard errors are clustered at the workplace level. Coefficients are multiplied by 100 so that they are converted to percentage point units.

Figure A.1: *Capital Contributions and Labor Income by Year*

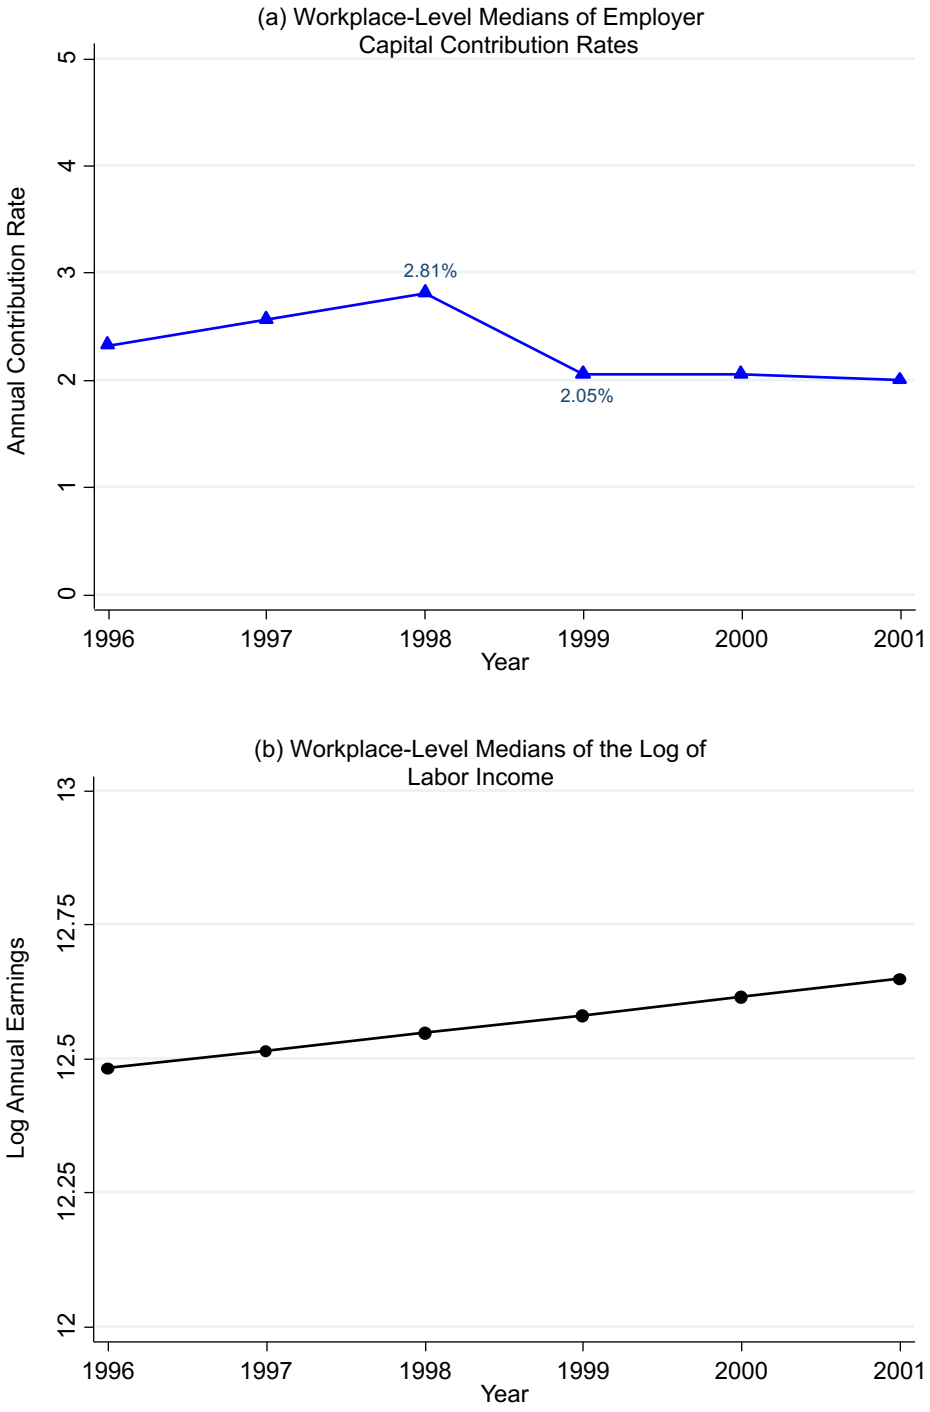

Notes: These figures plot means of different workplace-level outcomes by year, for years 1996-2001. The outcome in Panel A is the median contribution rate to employer-sponsored capital accounts, and the outcome in Panel B is the median of the log of annual labor income. A workplace is defined as the group of all employees with the same 2-digit occupation code in the same firm. The sample includes employees in private-sector firms with white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

Figure A.2: *Robustness of Employer Responses in Capital Contributions*

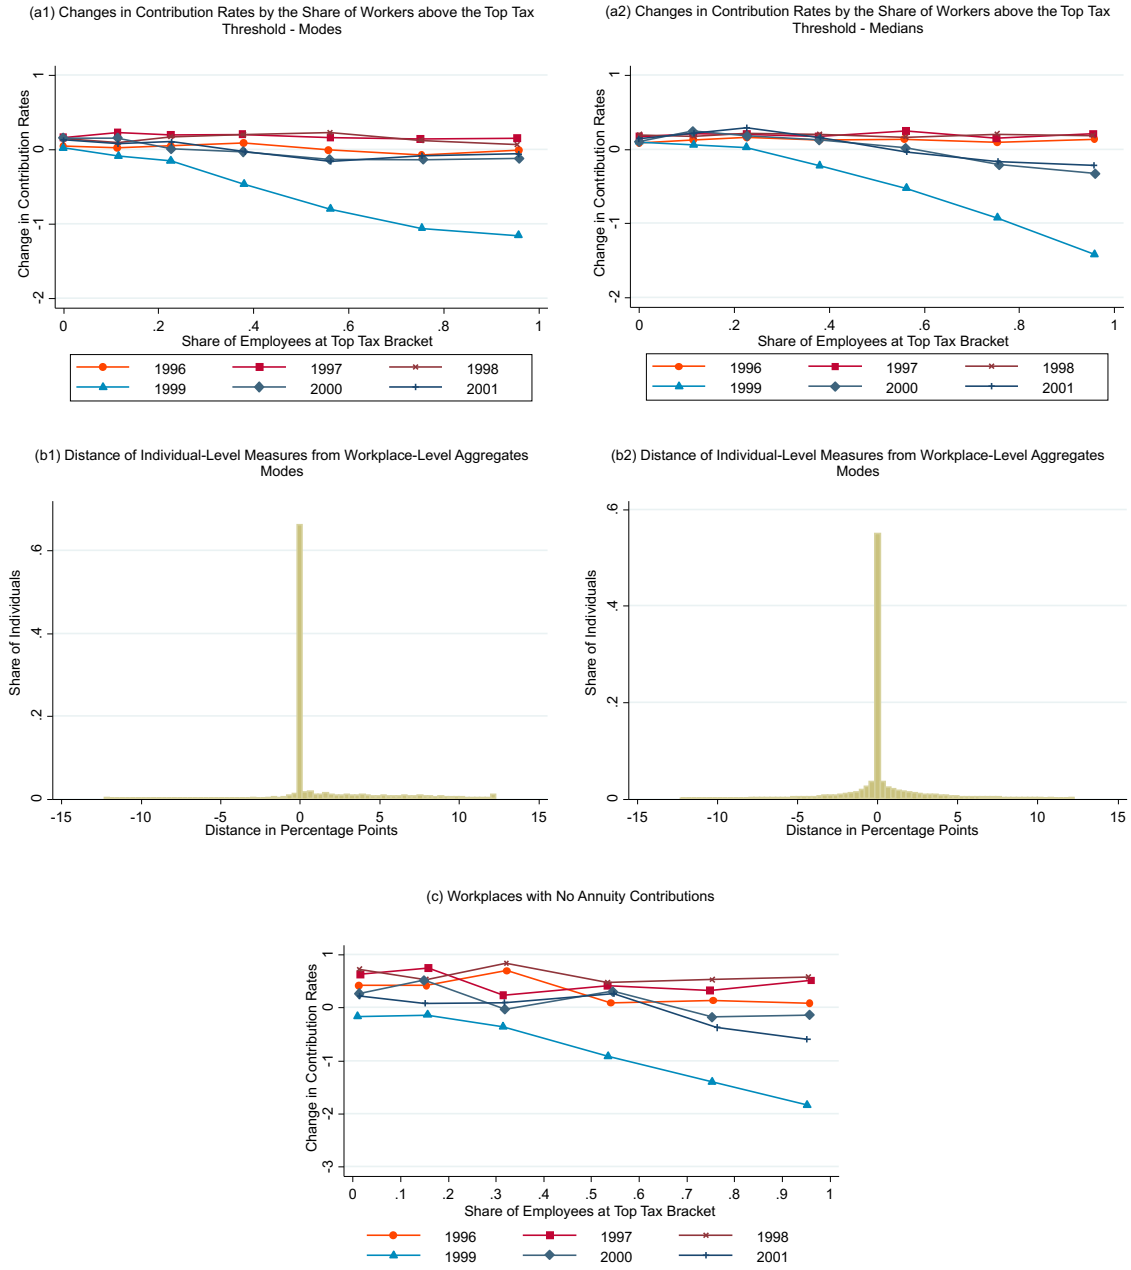

Notes: The sample for these figures includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees. Panel A.1 plots changes in workplace-level modes of contribution rates to capital accounts as a function of a workplace's share of employees above the top tax threshold, for years 1996-2001, including only workplaces with unique mode values. Panel A.2 plots workplace-level medians, including only the sample that is included in Panel A.1. These figures are plotted in the same way as the figure in Panel B of Figure 1.1. Panels B.1 and B.2 plot the distribution of the distance between employee-level capital contribution rates and workplace-level modes and medians, respectively. Panel C replicates Panel A.2 for all private-sector firms and white-collar occupations but constrains the sample to workplaces whose median annuity contribution rate in the years prior to the reform (1996-1998) was zero and whose median capital contribution rate in 1998 was larger than 1.5% (so that they had non-negligible potential to reduce capital contributions in response to the reform), with a total of 4,056 observations.

Figure A.3: *Employer Capital Contributions Responses by Size of the Workforce*

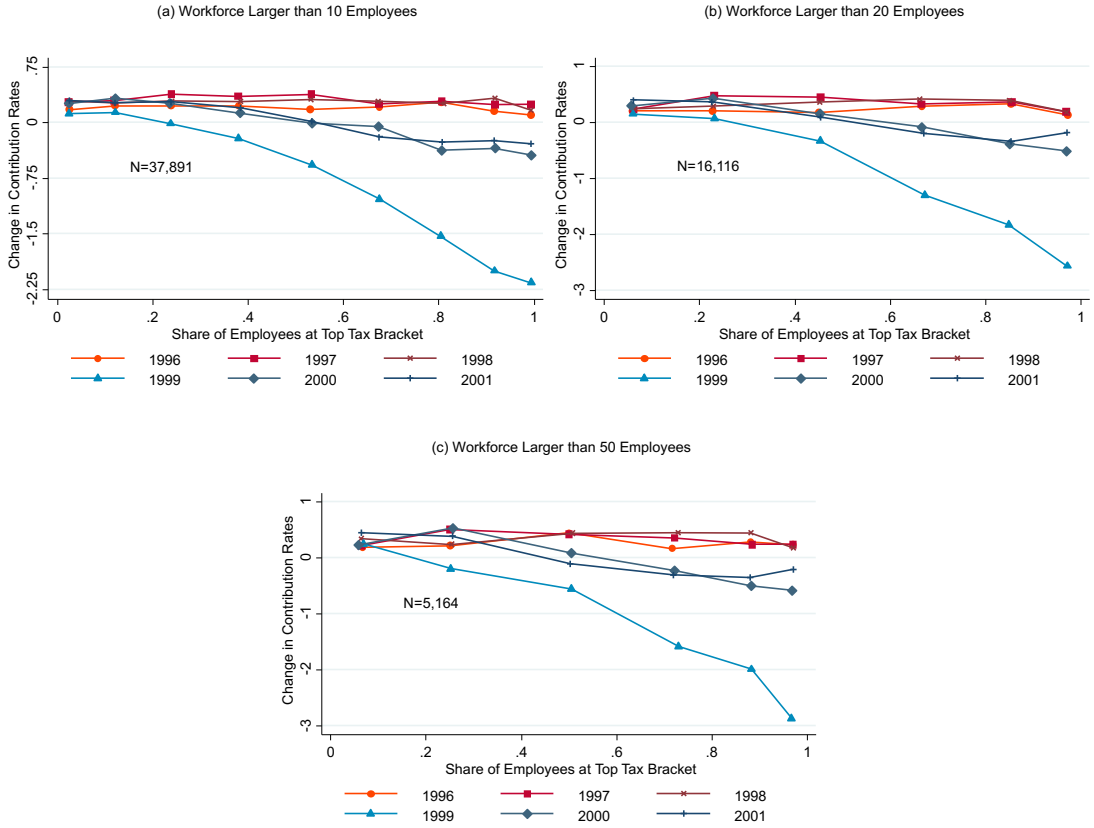

Notes: These figures plot changes in employers' contribution rates to capital pension accounts as a function of the share of their employees whose earnings placed them above the top labor-income tax threshold, for years 1996-2001. Panel A excludes workplaces with less than ten employees, Panel B excludes workplaces with less than twenty employees, and Panel C excludes workplaces with less than fifty employees. In each panel we indicate the number of included observations, denoted by N. The observation units are workplaces, defined as all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations and excludes self-employed individuals.

Figure A.4: *Changes in Employer vs. Individual Contributions to Capital Pension Accounts of Workers above the Top Tax Threshold*

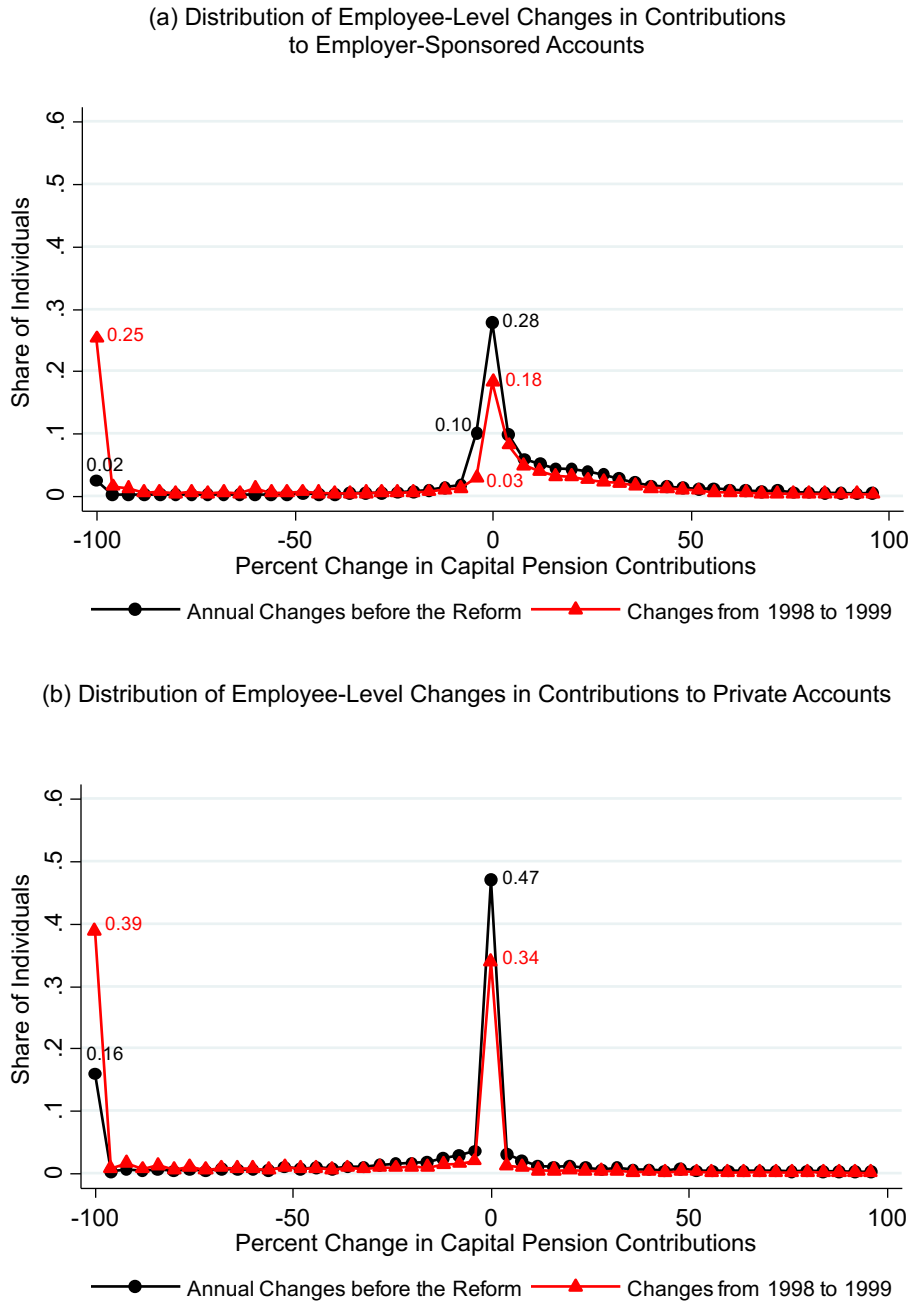

Notes: These figures plot the distribution of changes in employee-level capital pension contributions in percents relative to lagged contributions for individuals above the top tax cutoff for the years before the reform (black lines and circles) and from 1998 to 1999 (red lines and triangles). Panel A plots changes in contributions to employer-sponsored (401(k)-like) accounts, and Panel B plots changes in contributions to private (IRA-like) accounts. The figures include only individuals with positive lagged contributions to both types of accounts. Each point represents the floor of bins of 4% width, so that the point at 0% represents individuals with changes in the range [0%,4%). The curves for the distribution of annual changes in the years prior to the reform (in black lines and circles) include changes from 1996 to 1997 and 1997 to 1998. The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals, workplaces with less than five employees, and employees with earnings below the top income threshold.

Figure A.5: *Employer Capital Contributions and Labor Income by the Share of Workers above the Top Tax Threshold*

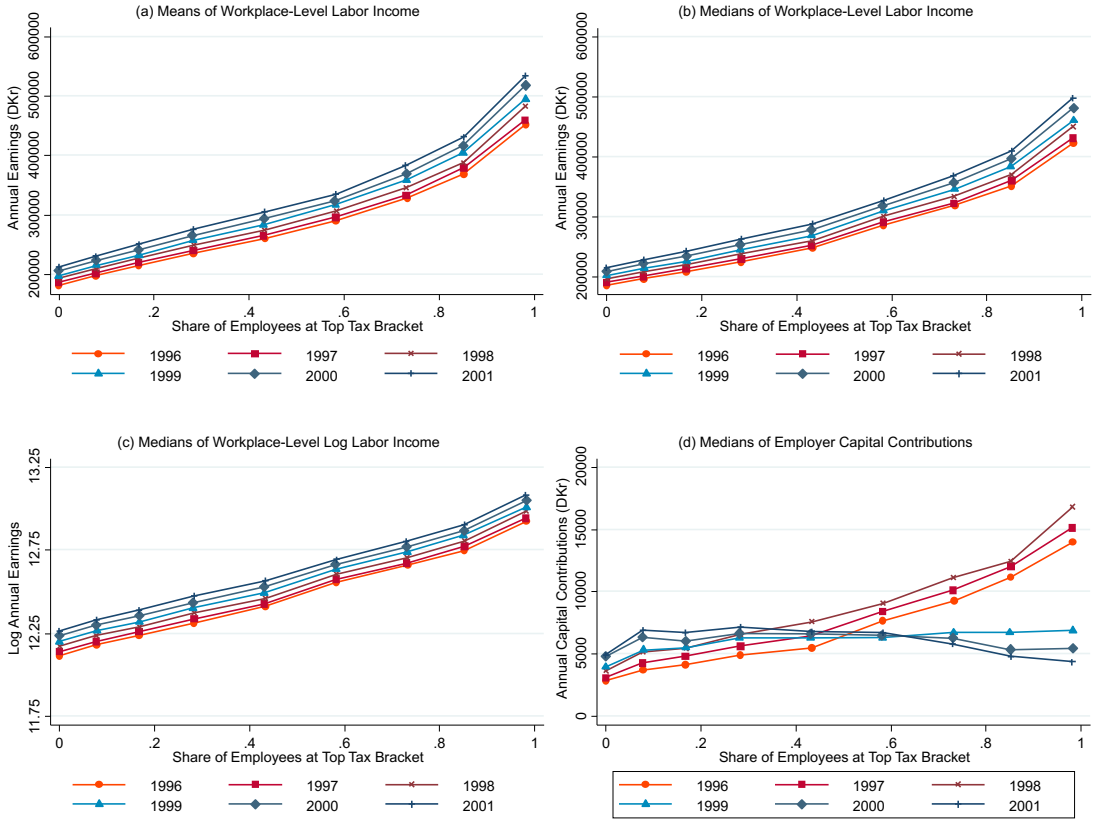

Notes: These figures plot different workplace-level outcomes as a function of the share of employees whose earnings placed them above the top labor-income tax threshold, for years 1996-2001. Panel A plots means of annual labor income, Panel B plots medians of annual labor income, Panel C plots medians of log annual labor income, and Panel D plots median employer contributions to capital pension accounts. The observation units are workplaces, defined as all employees with the same 2-digit occupation code in the same firm. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

Appendix B

Appendix to Chapter 2

B.1 Additional Institutional and Data Details

Copays for Prescription Drugs: The subsidy system has changed overtime, but has remained fairly generous over the time period we study. Prior to 2000, individuals paid 50% of the cost if the condition was not life threatening, while they paid 25% of the cost if the condition was life threatening. After 2000, individuals paid full amount for annual expenses up to 865 DKK ($133), then 50% for additional expenses in the range of DKK 865-1,410 ($133-216); 25% in the range of 1,510-3,045 ($216-468) DKK, and 15% for expenses over DKK 3,045 $468). After 2000, municipalities also gave various subsidies for drugs based on welfare status and income level. Appendix Figure A plots the average copay (in 2015 dollars) individuals paid for one pick up of each type of drug by year. While there are some fluctuations across time, and differences across drugs, individuals tend to pay about $10-$20 annually for these types of drugs. The reform seems to have the largest effect on the copay of anti-depressants, which increases from an approximately $18/pick up in 1999 to $26/pick in 2000. After 2000, the anti-depressant copay falls. Appendix Figure 2B plots the average total cost for the four types of drugs. It shows that the price of anti-depressants decreases substantially from 2000 onwards.

Physician Monitoring: General Practitioners are monitored by the Danish Patient Safety Authority. They make routine monitoring every third year for all GPs. In 2014, the Danish Patient Safety Authority initiated 244 cases against physicians due to various reasons - including the prescription behavior of physicians, inappropriate physician behavior, or breach of confidentiality.

However, only 2-3 cases are annually taken to court and most often these are due to breach of confidentiality cases.

Identification of Physicians: GP clinics are organized self-employed businesses, and operate with a provider number, which is fixed to an address. Regional councils supply provider numbers to geographical areas dependent on population density. Therefore, a physician cannot move their practice or provider number to another area to meet a specific demand. We identify physicians by their provider number (ydernumre). There are approximately 3500 "capacities" per year (one capacity serves app. 1500 patients), distributed between approximately 2500 provider numbers. There are more capacities than provider numbers because some providers employ one or more additional physicians. Hence, in most cases the provider number captures one specific GP, but in some cases it covers more physicians.

Labor Force Participation: We measure labor force participation as whether someone earned any labor income or self employment income within the calendar year.

Labor Income: While labor income rank is the primary labor income measure we use, we also test robustness using log real labor income, where we take the log of labor and self-employed earnings converted into 2015 dollars and then take the natural log plus one to include individuals who do not work.

Social Disability Insurance Receipt: We define social disability receipt by whether an individual receives social disability insurances within a year, defined for individuals less than 65. This is because the system switches to an old age pension after 65 for individuals born after July 1, 1939, and at age 67 for individuals born before July 1, 1939. To guarantee we aren't measuring pension income, we use age 65 to be conservative.

Sick Leave: Prior to April 2, 2007, private sector employers were obliged to pay at least weeks of sick pay, after which the municipality pays up to a year. From April 2, 2007 to June 2, 2008 private sector employees were obliged to pay up to 15 calendar days. From June 2, 2008-January 2, 2012, it became 21 calendar days, and after January 2, 2012 it became up to 30 calendar days. Therefore we create two measure of sick pay, the first is whether a private sector employee takes any sick pay (based on employer paid sick pay), and the other whether they take sick pay for more than at least two weeks (whether the municipality pays for sick pay). Since municipality sick pay includes absences due to parental leave, we set the measure of municipality sick pay equal to

missing if a person has had a child within the last two years. For both measures, we set them equal to missing for public sector workers.

Note that some common agreements in the private sector for wages during sickness beyond the required amount - therefore, these worker may not be in the municipality based sickness measure even if they have a sick leave absence longer than the required amount of time.

B.2 Theoretical Model

In this Appendix, we provide a one theoretical model that motivates the empirical identification strategy we use to estimate the effect that physician prescribing rates have on individual's prescription drug use and their labor supply. In particular, we show that when there are endogenous physician choices, and trends in health, under some assumptions which we detail below, we can identify the effect of physician prescribing rates exploiting a semi-exogenous separation of an individual from their doctor due to a cross municipality move. To identify the effect of the resulting change in physician prescribing rates from the move, first, we instrument for the change in physician prescription rates with the old physician's prescription rates to avoid the endogeneity from a concurrent health shock and the new physician choice. Second, we difference out the effects of a placebo group of individuals who don't move to allow flexibly for different health trends for individuals with previous doctors with different prescription drug rates.

Our empirical goal is to identify how physician prescribing rates affect individual prescription drug use and labor income and supply. We model prescription drug use and labor supply such that they are a function of observable characteristics, their doctors' prescription rate, their unobservable current health status, and an orthogonal random component.

$$D_{it} = \beta^D X_{it} + \gamma^D p_{j(it)} + \pi^D H_{it} + \nu_{it}^D$$

$$L_{it} = \beta^L X_{it} + \gamma^L p_{j(it)} + \pi^L H_{it} + \nu_{it}^L$$

We would like to identify γ^D and γ^L - the effect that their physician's prescribing rate has on their own drug use and their own labor supply respectively. However, since individuals may sort into physicians based on their health, there is likely a correlation between H_{it} and $p_{j(it)}$, which would bias estimates of $\hat{\gamma}^D$ and $\hat{\gamma}^L$. For example, when an individual hurts their back, they may switch to a GP who prescribes more opioids.

To understand this endogeneity problem, we model individual's health and how individuals may sort into physicians. We assume that individual's health, H_{it}, is the sum of a predictable function of past health, and a random component.

$$H_{it} = g(H_{it-1}, H_{it-2}, H_{it-3}, ...) + \epsilon_{it}$$

Individuals choose their doctor at time t as a function of their current health, their expected future health (which is just a function of their past health), and a random component that is orthogonal to their current and expected future health.

$$p_{j(it)} = f(H_{it}, H_{it-1}, H_{it-2}, H_{it-3}, ...) + \varepsilon_{it}$$

Individuals switch physicians if the benefit to do so, in terms of their expected future discounted utility, is higher than the cost of switching, C. Therefore, both in the cross-section and within an individual, differences in $p_{j(it)}$ are correlated with differences in health. In particular, if we think about switches of physicians within individuals: both the old physician's prescribing rate and the new physician's prescribing rate could potentially be correlated with the health shock. Within this framework, if someone leaves a low prescribing physician, it is because they have a negative health shock - similarly, if someone chooses a new doctor with a high prescribing rate, it is likely because they have had a negative health shock. To combat this endogeneity problem, our proposed identification strategy will focus on physician changes due to cross-municipality moves that we claim are unrelated to changes in health and their prior physician's prescribing rate.

Suppose an individual moves at time T to a new municipality. Prior to moving, they had physician k that they choose at some point, S, prior to T, $S < T$. The cross municipality move causes a sharp decrease in the cost of switching physicians, therefore, the probability individuals switch physicians at time T increases discretely. We claim that this increase in the probability of switching is unrelated to their physician's prescribing rate, causing a semi-exogenous separation at time T from the physician the individual choose at time S. Thus, if we were to compare the expected difference in the health of individual who moved at time T, to those who did not move at time T, it would *not* be a function of their prior physician's prescription drug rate: $\mathbb{E}(H_{iT}^M - H_{iT}^C | p_{j(iT-1)}) \neq h(p_{j(iT-1)})$.

The individuals who move are more likely to experience a change in their physician prescribing rates due to the separation with their old physician. Due to the random component in choosing a physician, there is mean reversion in the new physician choice, which means that the the change

in physician prescribing rates is inversely related to individual's prior physician prescribing rate. This is important, because we do not assume that individuals who move and individuals who don't move have the *same health* ($\mathbb{E}(H_{iT}^M - H_{iT}^C | p_{j(iT-1)}) = 0$), we cannot only compare movers to non-movers, but we also have to compare movers with different prior physician prescribing rates and thus different changes in physician prescribing rates.

Individuals who don't move still have some probability of separation, so the relative treatment between the groups is the difference in expected physician prescription changes based on the prior physician's prescription rate.

Therefore our identification strategy becomes:

$$\frac{(D_{iA}^M - D_{iB}^M) - (D_{iA}^C - D_{iB}^C)}{\mathbb{E}(p_{j(iA)}^M - p_{j(iB)}^M | p_{j(iB)}^M) - \mathbb{E}(p_{j(iA)}^C - p_{j(iB)}^C | p_{j(iB)}^C)}$$

In summary, this theoretical framework leads to the following proposed identification strategy, which has three components: first, we consider individuals who have a change in physician prescribing rates due to a between municipality move, second, we instrument for the change in physician prescription rates with the old physician's prescription rates, and third, we difference out the effects of a placebo group of individuals who don't move to allow flexibly for different trends for individuals with different previous doctors.

This identification procedure is based on the following assumptions:

1. Previous doctor's prescribing rate is predictive of change in prescription rate: $Cov(p_{j(iA)} - p_{j(iB)}, p_{j(iB)}) \neq 0$.

2. Other determinants of the relative change in prescription drug use and labor force participation between movers and non-movers are unrelated to the origin physician's prescribing rate, outside of its effects on the change in physician prescribing rates: $Cov(\epsilon_i, p_{j(iB)} | \nu_{it}^D, H_{it}^D) = 0$. There are multiple reasons this could fail:

   (a) $\mathbb{E}(H_{iT}^M - H_{iT}^C | p_{j(iT-1)}) = h(p_{j(iT-1)})$ - the difference in the mover and non-movers health at time T is related to the prior prescribing rate.

   (b) Physicians do have other characteristics: $\chi_{j(it)}$ that are correlated with $p_{j(it)}$ and with ν_{it}^D and ν_{it}^L.

In the paper, we address these potential threats to identification.

B.3 Moving Example

To understand the identification strategy we use to identify the effects of physician prescribing rates on individual outcomes, consider the following example: there are two individuals who move from Aarhus to Copenhagen. Before their move, one of them, A, visits a doctor in Aarhus who has a relatively high opioid prescribing rate - 10%, and another, B, visits a doctor that has a relatively low prescribing rate - 6%. After their move to Copenhagen, they are both separated from their doctors.

Due to some mean reversion in their doctor choice, the doctors that they choose in Copenhagen have a prescribing rate that is closer to the mean.[1] This means that on average they have different changes in the prescribing rates of their doctors after the move. In particular, A will have on average a decrease in their physician's prescription rate of opioids, while B will have on average an increase in their physician's prescription rate. Therefore, to understand the implications of a change in doctor prescribing rates, we can compare the change in prescription drug use and labor supply for these two individuals who have different physician prescribing rates prior to their move.

We don't compare individuals based on their actualized change physician prescribing rates before the new physician individuals choose may be based on their concurrent health. For example, consider the possibility that B hurts his back during the move, so picks a new doctor that has high prescribing rates. while A does not hurt his back. Thus the change in their physician prescribing rates will be endogenous to their changes in health as well. Therefore, we instrument for the change in physician prescribing rates with the pre-move physician's prescribing.

However, we may be concerned that there are other differences between these two individuals that are related to the prescribing rate of their different doctors, which leads them to have different trends in their drug use and labor supply. We therefore consider two other individuals who live in Aarhus: C, who has the same physician in Aarhus as A, and D, who has the same physician in Aarhus as B. However, unlike A and B, C and D stay in Aarhus. These two individuals are similar to the first, but their physician's prescribing rates do not change. We therefore can take the triple

[1] This doesn't necessarily have to be true, however, in Section 3.3.2, we show this empirically. On average there is a .75 reversion to the mean in terms of the new doctor's prescribing rates for all drugs.

difference in outcomes between the movers and non-movers ($A - C$ and $B - D$), the individuals with high prescribing physicians and low prescribing physicians ($(A - C) - (B - D)$), and the after minus before ($[(A_1 - C_1) - (B_1 - D_1)] - [(A_0 - C_0) - (B_0 - D_0)]$) to identify how different changes in prescribing rates affect individual's own usage and their labor supply.

B.4 Linking Patients to Physicians

We link patients to their primary care physician base on on the General Practitioner (GP) they saw most in the surrounding 3 years $(T-1, T, T+1)$.[2] We define "most" based on the number of different years they saw the GP (1, 2, or 3), and then break ties based on the total number of services charged to the GP. If there are still ties, we then choose a GP randomly from the tied GPs.

After we have assigned individuals to GPs, we drop individuals who are assigned to GPs with fewer than 1000 assigned patients, or who have patients that are more than 13% of the municipalities total population. The first is to ensure precision in our estimates of physician prescribing rates, and the second is to make sure that the physician is not an overwhelming share of the municipality market, which would make it difficult to adequately control for municipality-wide effects.

[2]Before we match individuals, we first drop all general practitioners who see fewer than 2000 patients ever, and 400 patients within a year to ensure the GP is in practice throughout the year and sufficiently involved in the health market.

B.5 Non-Parametric Specification

The specification we choose to estimate the effect of a physician prescribing rates on individual's own prescription drug use assumed that the effect was linear. Given that prescription drug use is a binary outcome, it is possible that a non-linear specification (e.g. probit or logit) would fit better. Additionally, if a prescription drug is addictive (e.g. opioids), then we might expect negative changes in the physician prescription rate to have a smaller effect than positive changes. To test this, we estimate the relationship between individual's change in drug use and the instrumented change in the physician prescribing rate non-parametrically by first binning the instrumented change in prescription drug rates into 20 equal size bins. For each bin B that spans (a, b), we estimate the effect of being a mover after the move on drug use (θ_B) for all individuals in that bin:

$$Drug_{it} = \theta_B After \times Mover_i + \Psi Mover_i + \mu_{o,d,A} + \epsilon_{it} \forall i : \hat{\Delta}_{p(i)} \in B_{(a.b)} \tag{B.1}$$

We run the regression on the observations three years prior to the move and the three years after the move, not including the year of the move. In Appendix Figure 5, we plot $\hat{\theta}_B$ by the mean value of $\hat{\Delta}_{p(i)}$ within the bin B, which non-parametrically characterizes the relationship between individual's drug use and the instrumented change in the physician's prescription drug rate.

Panel A plots the coefficients for opioid drug use and shows there is no evidence that the effect is non-linear. This suggests that the linear specification is sufficient. Panel B-D depicts the results for anti-inflammatories, anti-anxieties, and anti-depressants respectively. None show evidence of a non-linear effect.

B.6 Heterogeneity in Treatment Effects on Labor Income by Individual Outcomes

To understand the distributional consequences for the variation in physician prescribing rates, we analyze whether there are heterogeneous effects on the labor outcomes based on various individual characteristics. Just like Section 4.4, we look for heterogenous effects by age, gender, education, and previous occupation.

To estimate heterogenous effects by each characteristic, X_i, we first reestimate $\hat{\Delta}_{p(i)}$ as a function of X_i. We then estimate the following equation:

$$\begin{aligned} LaborIncome_{it} =& \theta_1 After \times \hat{\Delta}_{p(i)i} \times Mover_i \times X_i + \Gamma_1 After \times \hat{\Delta}_{p(i)i} \times X_i + \\ & \Psi_1 After \times Mover \times X_i + \beta_1 Mover \times X_i + \lambda_1 \hat{\Delta}_{p(i)i} \times X_i + \\ & \theta_0 After \times \hat{\Delta}_{p(i)i} \times Mover_i + \Gamma_0 After \times \hat{\Delta}_{p(i)i} + \\ & \Psi_0 After \times Mover + \beta_0 Mover + \lambda_0 \hat{\Delta}_{p(i)i} + \mu_{o,d,r,X_i} + \epsilon_{it} \end{aligned} \tag{B.2}$$

We include controls for differential effects of the move by age, gender, education, and previous income, and additionally interact this with X_i. We also include origin by destination by after and characteristics fixed effects.

Appendix Table 5 reports the coefficients θ_0, the effect of physician prescribing rates on labor income for those *without* the characteristic, and θ_1, the difference in the effect of physician prescribing rates between those with the characteristic and those without the characteristic. Column (1) reports heterogeneity by age ($Age > 42$), Column (2) reports heterogeneity by gender ($X_i = female$), Column (3) reports heterogeneity by education ($X_i = YearsEduc > 13$), and Column (4) reports heterogeneity by occupation in year $T-4$ ($X_i = BlueCollar$).

Panel A reports the coefficients for physician prescribing rates of opioids. There are no statistically significant differences by the individual characteristics we look at, though the standard errors are too large to rule out meaningful differences. Panels B-D present the results for the other prescription drugs we look at. For anti-inflammatories in Panel B, we don't find any significant differences based on characteristics. For anti-anxieties in Panel C, we see that older individuals have a larger positive effect than younger individuals, such that a one percentage point increase in prescribing rate of anti-anxieties leads to a .14 (se=.06) percentile increase in labor income

rank. Additionally, individuals who are highly educated have a smaller effect of the anti-anxiety prescribing rate than individuals with lower education, such that they have essentially zero effect of the anti-anxiety prescribing rate, while those with low education have a positive effect of .15 percentiles (se=.06) for a 1 percentage point increase the anti-anxiety prescribing rate. For anti-depressants in Panel D, we see no statistically significantly different effects by the individual characteristics we look at, though again, the standard errors are too large to rule meaningful sized effects out.

B.7 Figures and Tables

Figure B.1: *Municipality Map of Denmark*

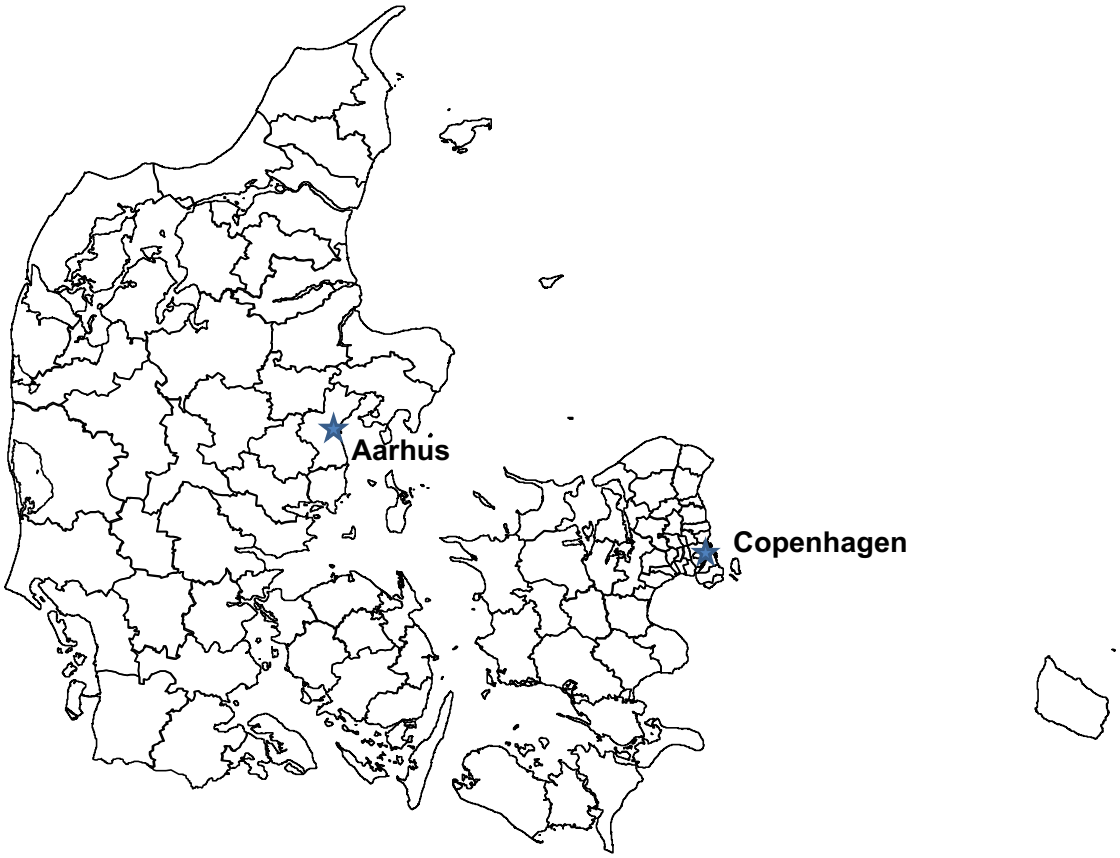

Notes: This figure shows a map of Denmark's 99 municipalities.

Figure B.2: *Price of One Pick-up for Prescription Drugs*

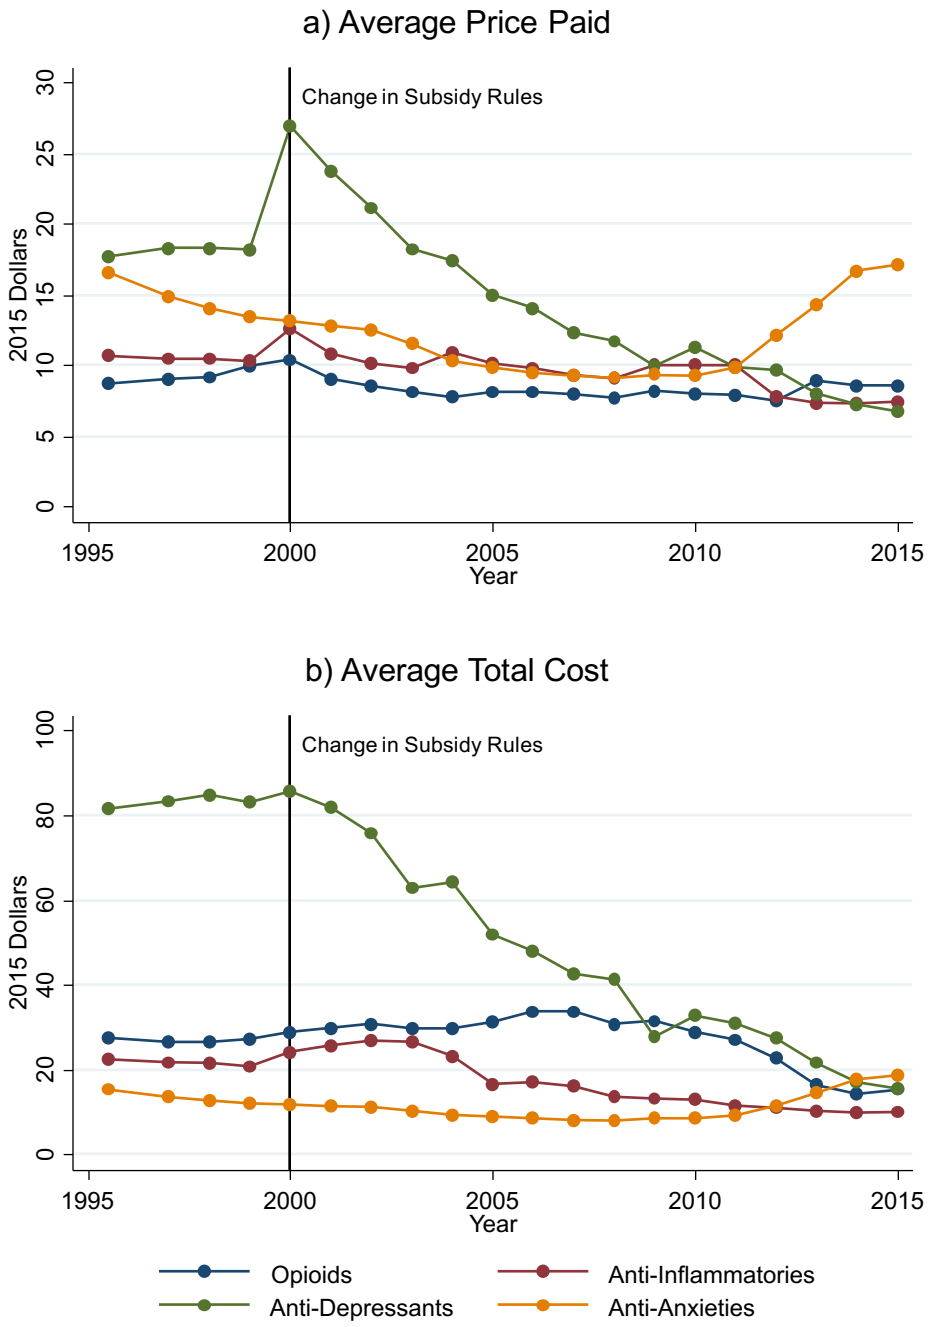

Notes: Panel A plots the average subsidized price that individuals paid in 2015 United States dollars for one pick-up in each year for each of the four types or prescription drugs we study: Opioids, Anti-Inflammatories, Anti-Anxieties, and Anti-Depressants. Panel B plots the non-subsidized cost of the prescription drugs by year for each of the four types of prescription drugs we study. In 2000, the government changed the subsidy system. Appendix A gives details for the pre and post 2000 subsidy systems.

Figure B.3: *Comparison of United States and Danish Prescription Opioid Drug Use*

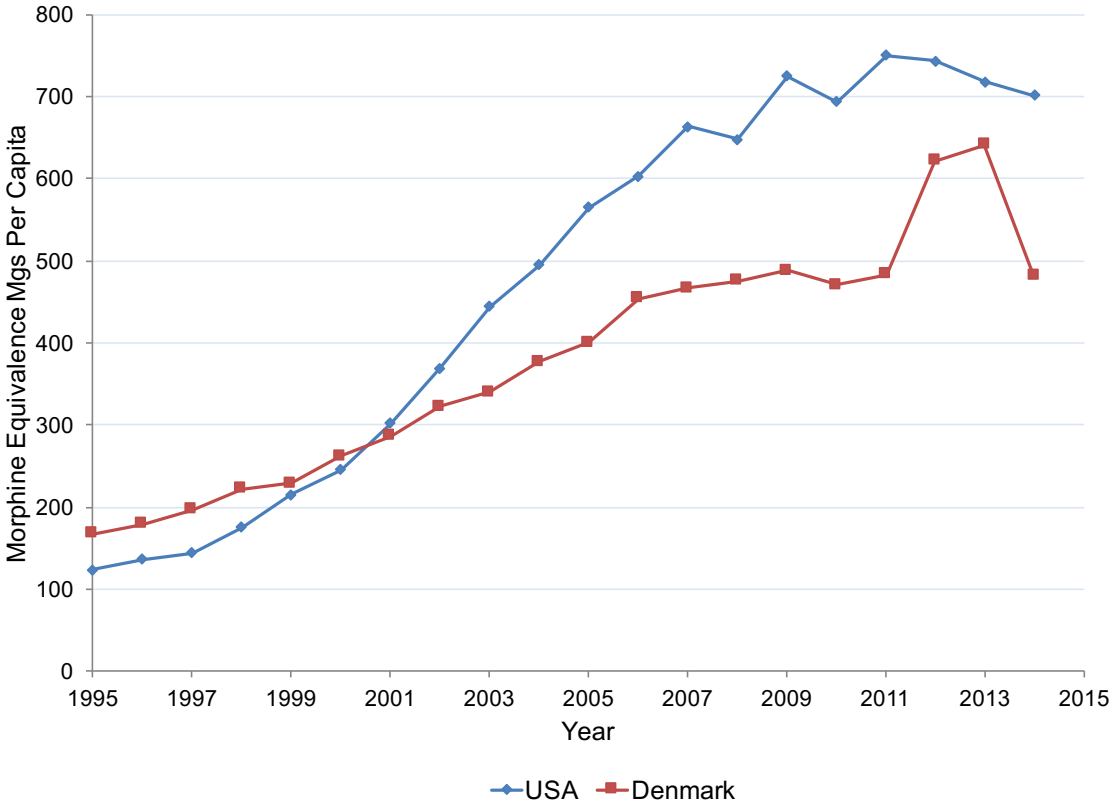

Source: Pain and Policy Studies Group at the University of Wisconsin-Madison, using International Narcotics Control Board and World Health Organization Population data

Notes: This figure uses data aggregated by the Pain and Policy Studies Group at the University of Wisconsin-Madison using data from International Narcotics Control Board and World Health Organization Population data. It graphs the per capita consumption of opioids for the United States (blue diamonds) and Denmark (red squares) in terms of Morphine Equivalence Milligrams from 1995-2014. The Pain and Policy Studies Groups developed the Morphine Equivalence (ME) metric using conversion factors from the WHO Collaborating Center for Drugs Statistics Methodology for the 6 principal opioids used to treat moderate to severe pain: Fentanyl, Hydromorphone, Methadone, Morphine, Oxycodone, and Pethidine.

Figure B.4: *Comparison of United States and Danish Prescription Drug Use - Non-Opioids*

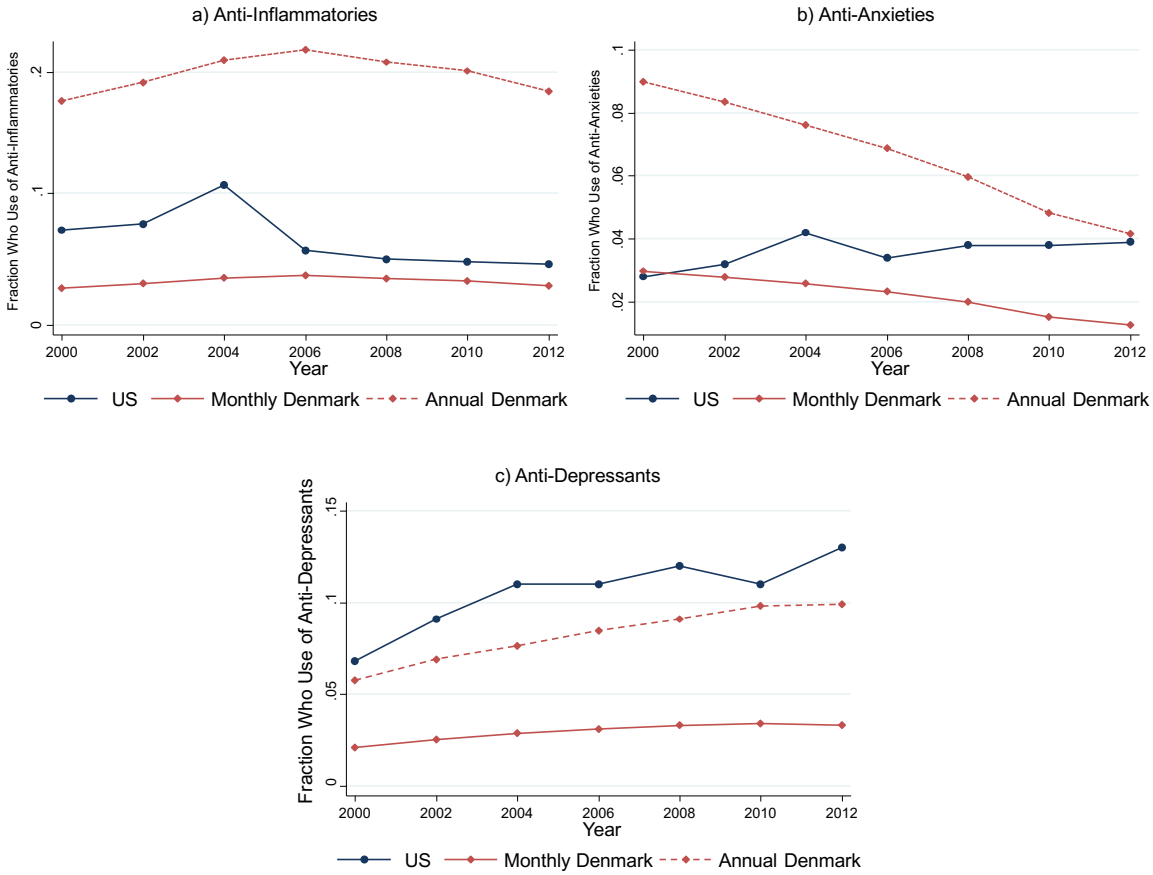

Notes: This figure plots estimates of average United States prescription drug use (dark blue circles) of from Kantor et al. (2015)'s analysis of the National Health and Nutrition Examination Survey (NHANES). This survey consisted of with seven cycles from 1999-2000 to 2011-2012. The metric reported is self-reported use of particular prescription drugs in the past 30 days. We compare this to the administrative Denmark data of indicators for any annual (dotted light red line with diamonds) or monthly pick ups (light red line with diamonds) of prescription drugs for 1999-2012 for individuals aged 32-70.

Figure B.5: *Non-Parametric Effect of a Change in Physician Prescribing Rates on the Change in Individual Drug Use*

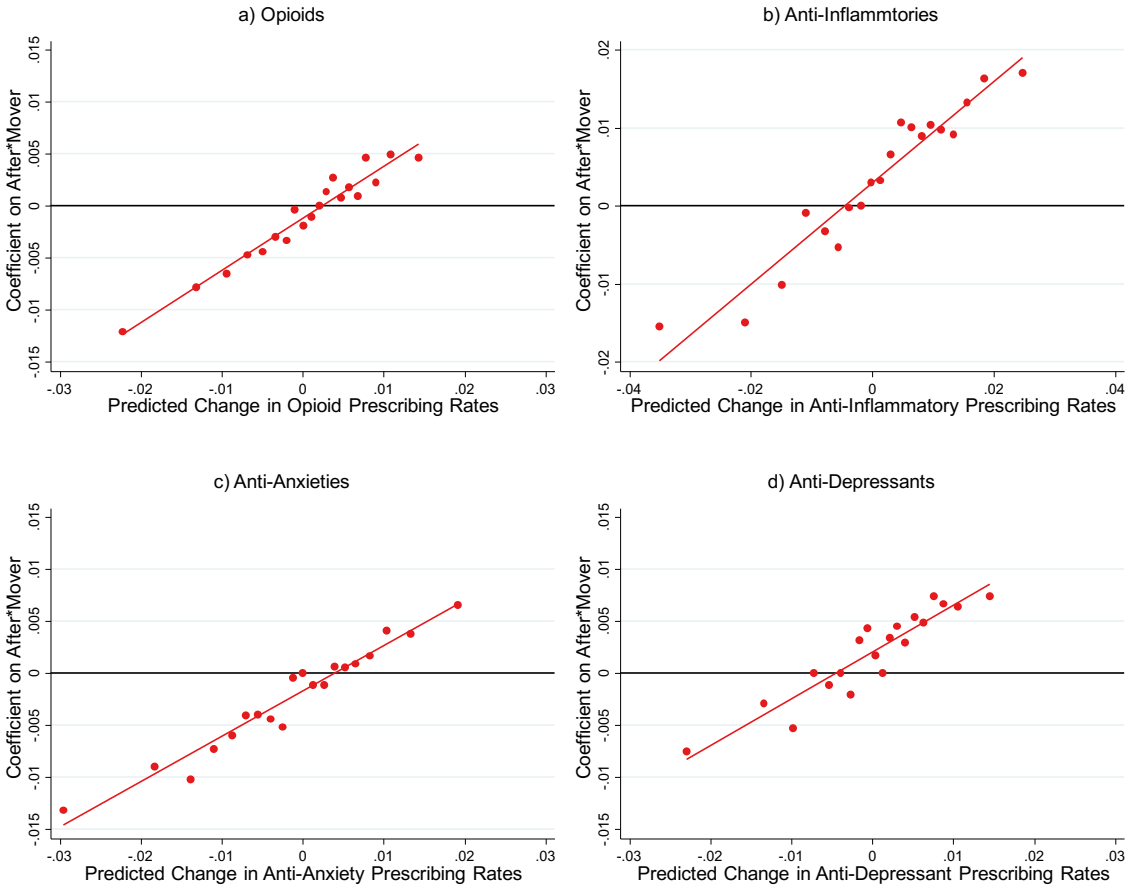

Notes: This figure shows the non-parametric relationship between the change in individual drug use and a change in the physician prescribing rates. This figure plots the coefficient in a regression of individual prescription drug use on an after indicator interacted with treatment indicator for different values of the predicted relative change in physician prescribing rates (θ_B from Equation 14 in Appendix E). Specifically, we bin the predicted relative change in prescribing rates into 20 equal sized bins. We then create indicators for each bin and interact them with an after indicator and a treatment indicator. We plot the coefficient on each interaction by the mean value of the relative predicted change in physician prescribing rates. Panel A plots the coefficients for opioid prescription drug use and opioid prescribing rates. Panel B plots the coefficients for anti-inflammatory prescription drug use and anti-inflammatory prescribing rates. Panel C plots the coefficients for anti-anxiety prescription drug use and anti-anxiety prescribing rates, and Panel D plots the coefficients for anti-depressant prescription drug use an anti-depressant prescribing rates.

Figure B.6: *Heterogeneity in the Effect of Physician Anti-Inflammatory Prescribing Rates by Distance of Move*

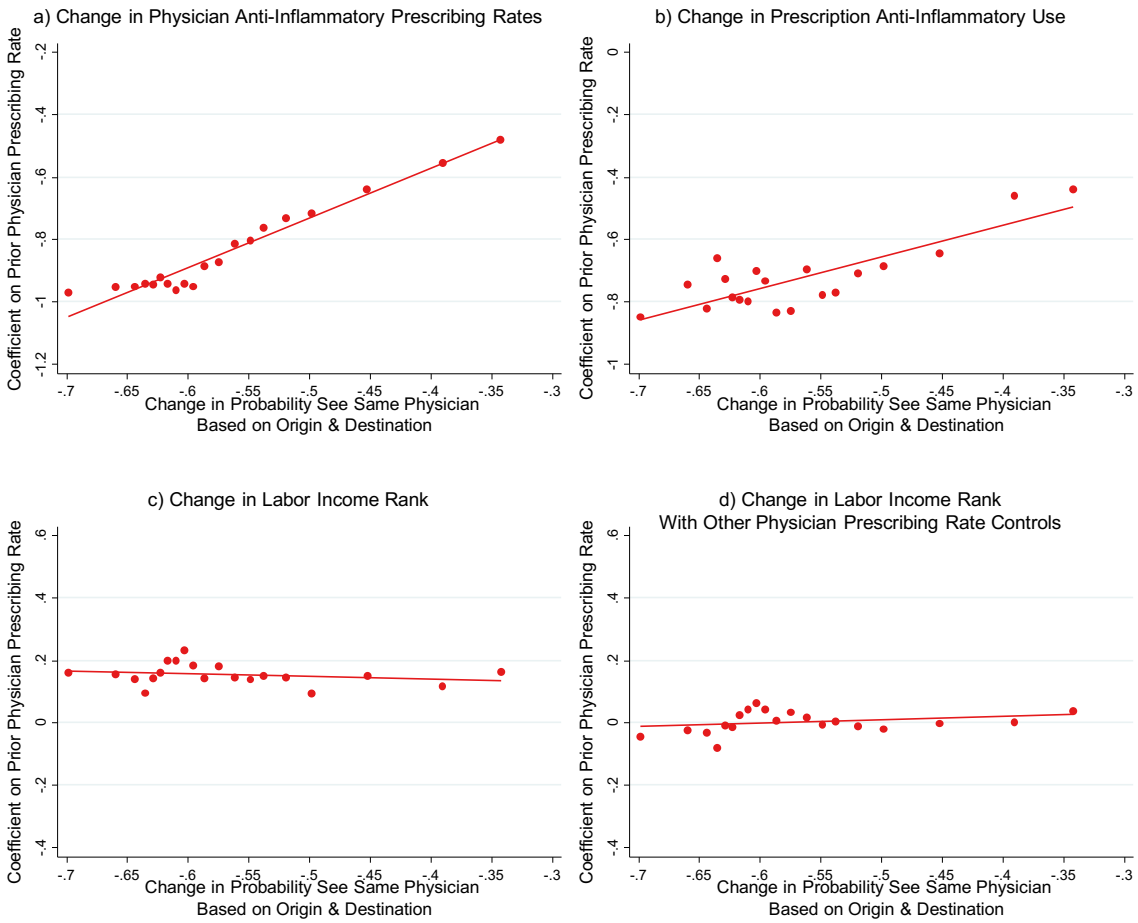

Notes: This figure shows heterogeneity in the effect of physician anti-inflammatory rates by the distance of the move - measured by the average separation rate with pre-move physicians. It looks at the effects of this heterogeneity on: the relationship between the change in physician prescribing rates and pre-move physician prescribing rates (Panel A), the relationship between the change in individual drug use and pre-move physician prescribing rates (Panel B), the relationship between the change in labor income rank and the pre-move physician's prescribing rates (Panel C and D). It replicates Figure 11, except for instead of opioid drug use and opioid prescribing rates, it looks at anti-inflammatory use and anti-inflammatory prescribing rates. See Figure 11 notes for additional details.

Figure B.7: *Heterogeneity in the Effect of Physician Anti-Anxiety Prescribing Rates by the Distance of Move*

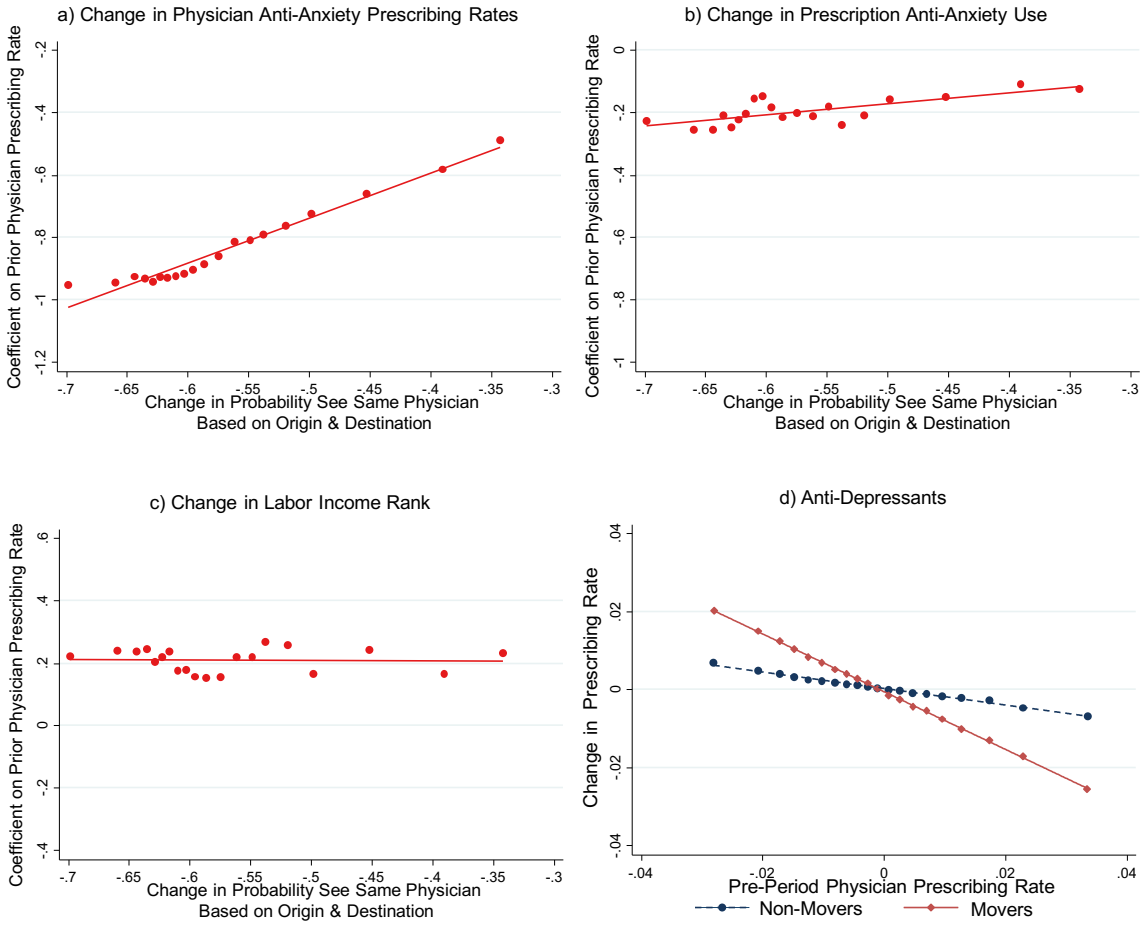

Notes: This figure shows heterogeneity in the effect of physician anti-anxiety rates by the distance of the move - measured by the average separation rate with pre-move physicians. It looks at the effects of this heterogeneity on: the relationship between the change in physician prescribing rates and pre-move physician prescribing rates (Panel A), the relationship between the change in individual drug use and pre-move physician prescribing rates (Panel B), the relationship between the change in labor income rank and the pre-move physician's prescribing rates (Panel C and D). It replicates Figure 11, except for instead of opioid drug use and opioid prescribing rates, it looks at anti-anxiety use and anti-anxiety prescribing rates. See Figure 11 notes for additional details.

Figure B.8: *Heterogeneity in the Effect of Physician Anti-Depressant Prescribing Rates by the Distance of Move*

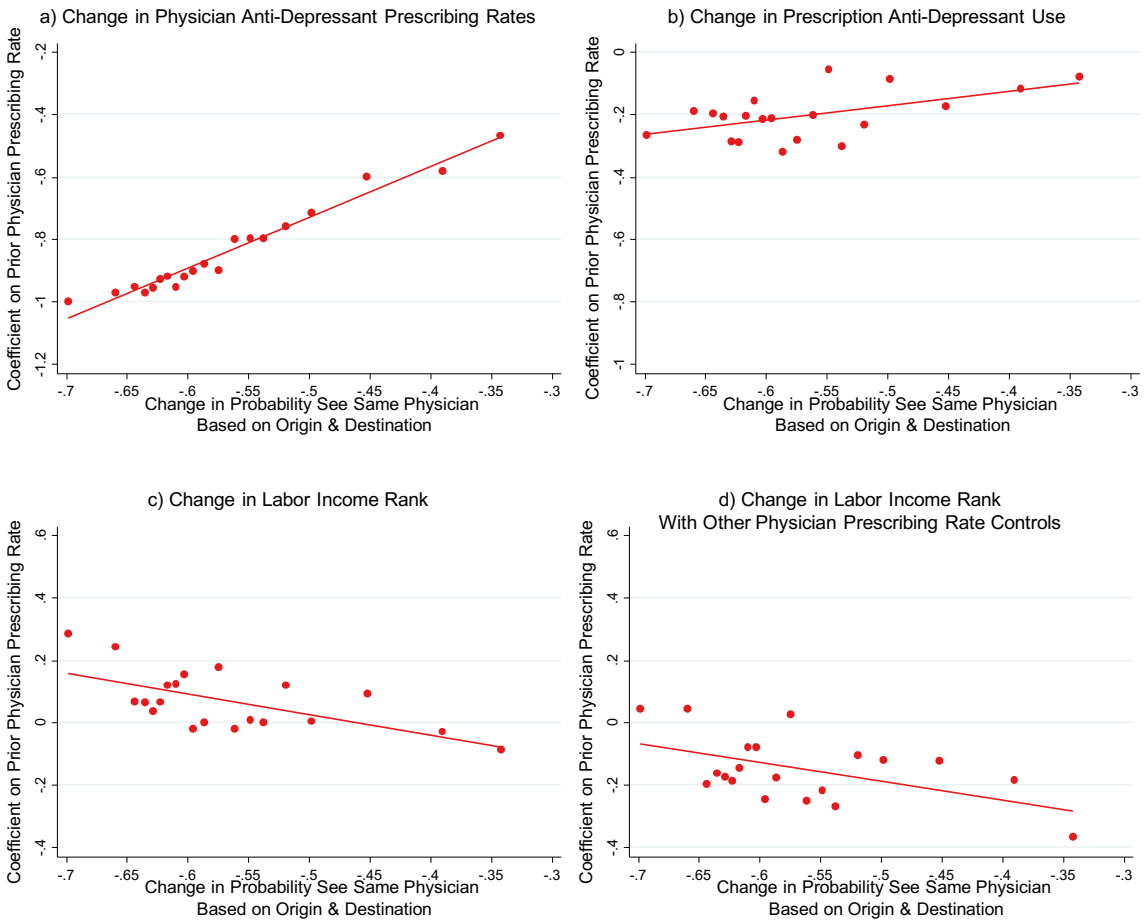

Notes: This figure shows heterogeneity in the effect of physician anti-depressant rates by the distance of the move - measured by the average separation rate with pre-move physicians. It looks at the effects of this heterogeneity on: the relationship between the change in physician prescribing rates and pre-move physician prescribing rates (Panel A), the relationship between the change in individual drug use and pre-move physician prescribing rates (Panel B), the relationship between the change in labor income rank and the pre-move physician's prescribing rates (Panel C and D). It replicates Figure 11, except for instead of opioid drug use and opioid prescribing rates, it looks at anti-depressant use and anti-depressant prescribing rates. See Figure 11 notes for additional details.

Table B.1: *Summary Statistics for Mover and Non-Mover Sample*

	Mover Sample		Non-Mover Sample	
	Mean (2)	SD (3)	Mean (5)	SD (6)
Individual Characteristics				
Year	2004	4.82	2004	4.83
Age	43	10.30	43	10.32
Year of Birth	1961	11.04	1961	11.06
Female	0.483	0.500	0.489	0.500
Yrs of Education	13.7	2.915	13.7	2.926
Rank of Ave Lab Inc T-8 to T-4	0.49	0.298	0.49	0.296
Blue Collar in T-4	0.430	0.495	0.456	0.498
Pre-Period Physician Prescribing Rates				
Opioids	-0.0004	0.018	-0.0002	0.018
Anti-Inflammatories	-0.0030	0.028	-0.0028	0.029
Anti-Anxieties	0.0041	0.024	0.0035	0.024
Anti-Depressants	0.0004	0.017	0.0004	0.017
Relative Predicted Change in Phy. Prescrib. Rates				
Opioids	0.0003	0.009	0.0002	0.009
Anti-Inflammatories	0.0008	0.014	0.0007	0.014
Anti-Anxieties	-0.0017	0.011	-0.0014	0.012
Anti-Depressants	-0.0002	0.009	-0.0002	0.010
Prescription Drug Use				
Opioids	0.069	0.253	0.062	0.241
Anti-Inflammatories	0.193	0.394	0.186	0.389
Anti-Anxieties	0.068	0.252	0.060	0.237
Anti-Depressants	0.091	0.287	0.074	0.262
Labor Market Outcomes				
Labor Income Rank	0.497	0.299	0.508	0.289
Labor Force Participation	0.811	0.392	0.825	0.380
Pos Labor Income Rank	0.580	0.266	0.585	0.254
Ln (Labor Income +1)	10.05	4.985	10.28	4.845
Employer Sick Pay	0.092	0.289	0.083	0.276
Municipality Sick Pay	0.095	0.293	0.082	0.274
Disability Reciept	0.072	0.259	0.076	0.265
N Observations	6,344,622		11,946,802	

Notes: This table reports the mean and standard deviation of the main set of variables we use in the paper for the mover sample (Columns (1)-(3)) and the non-mover control sample (Columns (4)-(6)) from three years prior to the "move" and up to three years after the "move". The non-mover control sample is matched on age, education, gender, quartiles of pre-period physician's prescribing rate, and quartile rank of average income from T-8 to T-4, as well as the year. Rank of average labor income from T-8 to T-4 is calculated within cohort, age, and year groups, where T refers to the year of the move. Pre-period physician prescribing rates are calculated by first variation out from individual's age, gender, the year, and their education and calculating physician effects. The relative predicted change in physician prescribing rates is calculated based on the difference between the mover and the non-mover control sample in the predicted change in physician prescribing rates based on the pre-period physician prescribing rates. Prescription drug use is an indicator for picking up a prescription a drug from the particular class within a year. Labor income rank is defined as rank of labor and self employed income within cohort, gender, and age groups. Labor force participation is an indicator for having positive labor or self employment income. Positive labor income rank is the labor income rank defined only for individuals with positive labor force participation. Employer sick pay is defined as whether an individual received any employer sick pay, which is defined only for workers in the private sector and individuals who have positive labor force participation. Municipality sick pay is defined as whether individuals received sick pay from the municipality, which kicks in after employers stop paying sick pay - which is generally at 2-4 weeks. It also includes paternity leave so it is only defined over the set of private sector workers who have positive labor force participation and who are not women who have had a baby in that year or the year previous. Disability receipt is defined for individuals less than 65.

Table B.2: *The Effect of Characteristics on Prescription Drug Use*

	Prescription Drug Use			
X:	Old	Female	Educated	Blue-Collar
	(1)	(2)	(3)	(4)
A. Opioids				
Constant	0.0438***	0.0557***	0.115***	0.0403***
	(0.0001)	(0.0001)	(0.0013)	(0.0001)
X	0.0418***	0.0120***	-0.0538***	0.0281***
	(0.0001)	(0.0001)	(0.0013)	(0.0001)
B. Anti-Inflammatories				
Constant	0.157***	0.169***	0.316***	0.153***
	(0.0001)	(0.0001)	(0.0021)	(0.0001)
X	0.0663***	0.0322***	-0.132***	0.0554***
	(0.0002)	(0.0002)	(0.0021)	(0.0002)
C. Anti-Anxieties				
Constant	0.0347***	0.0487***	0.0547***	0.0470***
	(0.0001)	(0.0001)	(0.0013)	(0.0001)
X	0.0619***	0.0251***	0.00617***	0.0129***
	(0.0001)	(0.0001)	(0.0013)	(0.0001)
D. Anti-Depressants				
Constant	0.0626***	0.0603***	0.127***	0.0608***
	(0.0001)	(0.0001)	(0.0014)	(0.0001)
X	0.0351***	0.0353***	-0.0497***	0.0119***
	(0.0001)	(0.0001)	(0.0014)	(0.0001)
N	15727935	15727935	15727935	12407329

Notes: This table runs a regression of prescription drug use on an indicator of individual characteristic. In Column (1) that characteristics is being older than 42, in Column (2) the characteristics is being female, in Column (3) that characteristics is having more than 14 years of education, and Column (4) that characteristics is having a blue collar occupation in T-4. Panel A reports the results when the outcome variable is opioid use with the year, Panel B reports the results when the outcome variable is anti-inflammatory use, Panel C reports the results when the outcome variable is anti-anxiety use, and Panel D reports the results when the outcome measure is anti-depressant use. These regressions are run on the same sample as Table 2.3.

Table B.3: *The Effects of Physician Prescribing Rates on Prescription Drugs Horse Race*

	Prescription Drug Use			
	Opioids (1)	Anti-Inflammatories (2)	Anti-Anxieties (3)	Anti-Depressants (4)
After × Mover × Predicted Δ in Physician Prescribing Rates	0.609*** (0.0513)	0.704*** (0.0382)	0.515*** (0.0478)	0.460*** (0.0382)
N	15,324,329	15,324,329	15,324,329	15,324,329

Notes: This table replicates Table 2.2 but includes the relative predicated change in prescribing rates for all drugs, so that the coefficients from four columns come from the same regression.

Table B.4: *The Effects of Physician Prescribing Rates on Other Labor Outcomes - Not Horse Race*

	Labor Income Rank (1)	LFP (2)	Log Labor Income (3)	Labor Inc. Rank Pos. LFP (4)	Sick Pay Employer (5)	Sick Pay Municipality (6)	DI Receipt (7)
A. Opioids							
After×Δ×Mover	-0.117***	-0.145**	-1.923**	-0.0732	-0.0152	0.0726	0.0176
	(0.0323)	(0.0490)	(0.591)	(0.0392)	(0.0563)	(0.0577)	(0.0335)
B. Anti Inflammatories							
After×Δ×Mover	-0.0354	-0.0454	-0.667	-0.00194	-0.0425	0.0238	0.0422*
	(0.0199)	(0.0295)	(0.356)	(0.0238)	(0.0340)	(0.0345)	(0.0197)
C. Anti Anxieties							
After×Δ×Mover	-0.0346	-0.0249	-0.416	-0.00791	-0.0288	0.00491	-0.0118
	(0.0248)	(0.0371)	(0.447)	(0.0298)	(0.0424)	(0.0432)	(0.0246)
D. Anti-Depressants							
After×Δ×Mover	-0.110***	-0.0858	-1.429*	-0.0902*	-0.0576	-0.0526	0.0517
	(0.0320)	(0.0480)	(0.579)	(0.0386)	(0.0548)	(0.0566)	(0.0337)
N	15171445	15171445	15171445	12469064	8964761	8587148	13750107

Notes: This table replicates Table 2.6, except instead of including the relative change of physician prescribing rates all together in the same regression, we estimate the coefficients in separate regressions, one for each panel within a column.

Table B.5: *Heterogeneity of the Effects on Labor Income Rank - Horse Race*

	Labor Income Rank			
X:	Old	Female	Educated	Blue-Collar
	(1)	(2)	(3)	(4)
A. Opioids				
After×Δ×Mover	-0.106	-0.166**	-0.189***	-0.191**
	(0.0613)	(0.0608)	(0.0555)	(0.0660)
After×Δ×Mover×X	-0.0186	0.0815	0.101	0.0993
	(0.0862)	(0.0866)	(0.0882)	(0.0960)
B. Anti-Inflammatories				
After×Δ×Mover	0.0224	0.00120	-0.0364	0.0460
	(0.0293)	(0.0293)	(0.0265)	(0.0312)
After×Δ×Mover×X	-0.0397	0.00335	0.0663	-0.0719
	(0.0411)	(0.0414)	(0.0419)	(0.0459)
C. Anti-Anxieties				
After×Δ×Mover	0.00841	0.0831	0.153***	0.0719
	(0.0449)	(0.0449)	(0.0415)	(0.0478)
After×Δ×Mover×X	0.141*	-0.0507	-0.155*	-0.0210
	(0.0637)	(0.0637)	(0.0649)	(0.0711)
D. Anti-Depressants				
After×Δ×Mover	-0.143**	-0.0544	-0.0617	-0.0983
	(0.0494)	(0.0505)	(0.0463)	(0.0544)
After×Δ×Mover×X	0.106	-0.0519	-0.0661	0.0521
	(0.0708)	(0.0710)	(0.0719)	(0.0804)
N	15171445	15171445	15171445	12096004

Notes: This table replicates Table 2.3, however, labor income rank is the outcome variable, and we include all of the physician prescribing rates interacted fully with an indicator for after, the predicted relative change in prescribing rates, and indicator for mover and an indicator for the characteristics. Therefore now each column reports coefficients from the same regression. We also include the controls from Table 4 Column (3) also interacted with the characteristic. We measure labor income as taxable labor and self-employed earnings. We convert labor income into percentile ranks within an individual's year of birth, the year, and their gender using the full sample of the Danish population (not just the movers and non-mover control sample). It is on a scale from 0 (lowest income) to 1 (highest income). Appendix Table 2.5 reports the coefficients when we separately control for each prescribing rate in separate regressions.

Table B.6: *Heterogeneity of the Effects on Labor Income Rank - Not Horse Race*

	Labor Income Rank			
X:	Old	Female	Educated	Blue-Collar
	(1)	(2)	(3)	(4)
A. Opioids				
After×Δ×Mover	-0.159***	-0.124**	-0.114**	-0.147**
	(0.0421)	(0.0420)	(0.0388)	(0.0461)
After×Δ×Mover×X	0.130*	0.0183	-0.0215	0.0556
	(0.0601)	(0.0604)	(0.0612)	(0.0673)
B. Anti Inflammatories				
After×Δ×Mover	-0.0361	-0.0314	-0.0559*	-0.00480
	(0.0260)	(0.0262)	(0.0238)	(0.0278)
After×Δ×Mover×X	0.0201	0.000469	0.0337	-0.0424
	(0.0367)	(0.0370)	(0.0374)	(0.0410)
C. Anti Anxieties				
After×Δ×Mover	-0.0873**	-0.0202	0.0201	-0.0372
	(0.0319)	(0.0323)	(0.0300)	(0.0344)
After×Δ×Mover×X	0.155***	-0.0240	-0.0976*	0.0144
	(0.0460)	(0.0459)	(0.0466)	(0.0515)
D. Anti-Depressants				
After×Δ×Mover	-0.183***	-0.0935*	-0.0873*	-0.123**
	(0.0412)	(0.0422)	(0.0389)	(0.0454)
After×Δ×Mover×X	0.162**	-0.0322	-0.0699	0.0411
	(0.0594)	(0.0595)	(0.0602)	(0.0675)
N	15171445	15171445	15171445	12096004

Notes: This table replicates Appendix Table 2.4 except it estimates the coefficients for each Panel in separate regressions, not simultaneously controlling for the four types of physician prescribing rates.

Appendix C

Appendix to Chapter 3

Figure C.1: *Distribution of Year of Birth of Sample*

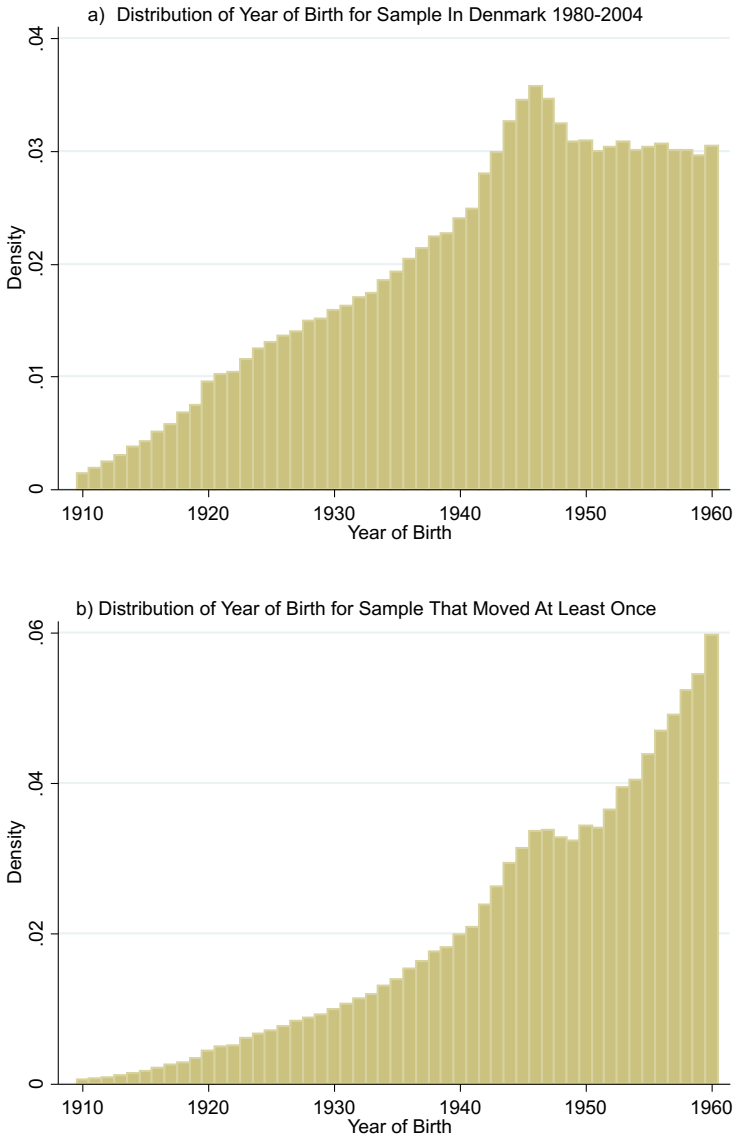

Notes: Appendix Figure 1a plots the distribution of the year of birth of the sample of individuals who lived in Denmark from 1980-2004. Append Figure 1b plots the distribution of the year of birth for the sample of individuals who moved at least once between municipalities between 1984-2004 and were present in Denmark from 1980-2004.

Figure C.2: *Distribution of Number Years Spent in Municipality*

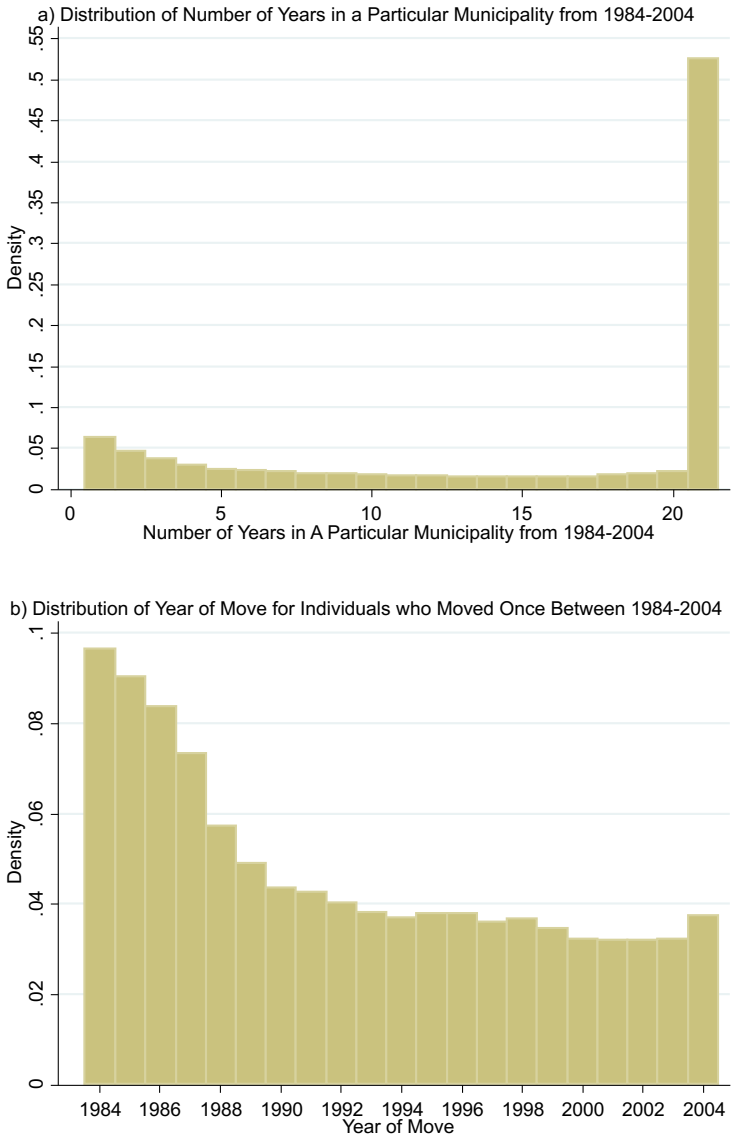

Notes: Appendix Figure 2a plots the distribution of the number of year an individual spent in a particular municipality from 1984-2004 for individuals from cohorts 1910-1960 and who were in Denmark from 1980-2004. Appendix Figure 2b plots the distribution of the year of move for the set of individuals who moved only once between 1984-2004.