

NBER WORKING PAPER SERIES

ADJUSTMENT COSTS, FIRM RESPONSES, AND MICRO VS. MACRO LABOR SUPPLY ELASTICITIES:
EVIDENCE FROM DANISH TAX RECORDS

Raj Chetty
John N. Friedman
Tore Olsen
Luigi Pistaferri

Working Paper 15617
<http://www.nber.org/papers/w15617>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2009

We would like to thank David Card, Stephen Coate, Edward Glaeser, James Hines, Han Hong, Lawrence Katz, Henrik Kleven, Claus Kreiner, Patrick Kline, Erzo Luttmer, Robert Moffitt, John Pencavel, Emmanuel Saez, Esben Schultz, and numerous seminar participants for helpful suggestions and valuable discussion. We are extremely grateful to Mette Ejrnæs and Bertel Schjerning at the Centre for Applied Microeconometrics at the University of Copenhagen, Frederik Hansen at the Ministry of Finance, Peter Elmer Lauritsen at Statistics Denmark, as well as Anders Frederiksen, Paul Bingley, and Niels Chr. Westergård-Nielsen at Aarhus Business School for help with the data and institutional background. Gregory Bruich, Jane Choi, and Keli Liu provided outstanding research assistance. We are also extremely grateful to Chris Walker and the Harvard Research Computing group for technical assistance. Support for this research was provided by the Robert Wood Johnson Foundation and NSF Grant SES-0645396. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Raj Chetty, John N. Friedman, Tore Olsen, and Luigi Pistaferri. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records

Raj Chetty, John N. Friedman, Tore Olsen, and Luigi Pistaferri

NBER Working Paper No. 15617

December 2009, Revised January 2011

JEL No. E62,H2,J22

ABSTRACT

We show that the effects of taxes on labor supply are shaped by interactions between adjustment costs for workers and hours constraints set by firms. We develop a model in which firms post job offers characterized by an hours requirement and workers pay search costs to find jobs. In this model, micro elasticities are smaller than macro elasticities because they do not account for adjustment costs and firm responses. We present evidence supporting three predictions of the model by analyzing bunching at kinks using the universe of tax records in Denmark. First, larger kinks generate larger taxable income elasticities because they are more likely to overcome search costs. Second, kinks that apply to a larger group of workers generate larger elasticities because they induce changes in hours constraints. Third, firms tailor job offers to match workers' aggregate tax preferences in equilibrium. Calibrating our model to match these empirical findings, we obtain a lower bound on the intensive-margin macro elasticity of 0.34, an order of magnitude larger than the estimates obtained using standard microeconomic methods for wage earners in our data.

Raj Chetty
Department of Economics
Harvard University
1805 Cambridge St.
Cambridge, MA 02138
and NBER
chetty@fas.harvard.edu

John N. Friedman
Harvard Kennedy School
Taubman 356
79 JFK St.
Cambridge, MA 02138
and NBER
jfriedman@post.harvard.edu

Tore Olsen
Harvard University
tolsen@fas.harvard.edu

Luigi Pistaferri
Department of Economics
579 Serra Mall
Stanford University
Stanford, CA 94305-6072
and NBER
pista@stanford.edu

I Introduction

The vast theoretical and empirical literature on taxation and labor supply generally assumes that workers can freely choose jobs that suit their preferences. This paper shows that the effect of taxes on labor supply is shaped by two factors that limit workers' ability to make optimal choices: adjustment costs and hours constraints endogenously set by firms. We present quasi-experimental evidence showing that these forces attenuate microeconomic estimates of labor supply elasticities and develop a method of recovering macro elasticities from micro estimates.

To structure our empirical analysis, we develop an model of intensive-margin labor supply with job search costs and endogenous hours constraints. We model hours constraints by assuming that each firm requires its employees to work a fixed number of hours because of an ex ante commitment to a production technology.¹ Workers draw offers from the aggregate distribution of hours offered by firms. Workers can search for jobs that require hours closer to their unconstrained optimum by paying search costs. In equilibrium, the number of jobs posted by firms at each level of hours equals the number of workers who select those hours after the search process is complete. The aggregate distribution of workers' preferences therefore determines the hours constraints imposed by firms in equilibrium.² However, most individuals do not work their unconstrained optimal number of hours because of search costs.

Our model produces a divergence between micro and macro labor supply elasticities. We define the macro elasticity as the effect of variation in taxes across economies on average hours of work. We show that the macro elasticity always equals the "structural" labor supply elasticity ε , the parameter of individuals' utility functions that determines elasticities absent frictions. In contrast, micro elasticities – defined as the effect of tax changes or kinks in non-linear tax systems that affect subgroups of workers – are attenuated relative to ε because of search costs and hours constraints.

The model generates three testable predictions about how search costs and hours constraints affect the labor supply (or taxable income) elasticities observed in micro studies. First, the observed elasticity increases with the size of the tax variation from which the estimate is identified. Intuitively, large tax changes prompt more individuals to pay search costs and find a new job.

¹We focus on hours constraints in our model for simplicity, but they should be interpreted more broadly as technological constraints on job characteristics (e.g. training, effort, benefit packages). When jobs have multiple characteristics beyond hours, the key predictions of our model apply to the taxable income elasticity rather than the hours elasticity.

²This endogenous determination of wage-hours offers is the key innovation in this model relative to the few existing models of hours constraints, in which firms' technologies fully determine the distribution of wage-hours packages (e.g. Rosen 1976).

Analogously, larger kinks induce more individuals to pay search costs to find a job that places them at the kink. Second, the observed elasticity increases with the number of workers affected by a tax change or kink. Changes in taxes induce changes in labor supply not just by making individuals search for different jobs, but also by changing the distribution of hours offered by firms in equilibrium. Because changes in taxes that affect a larger group of individuals induce larger changes in hours constraints, they generate larger observed elasticities. Furthermore, tax changes may affect even the labor supply of workers whose personal tax incentives are *unchanged* by distorting their coworkers' incentives and inducing changes in hours constraints. Finally, the model predicts a correlation between firm responses and individual responses to taxes. Because firms cater to workers' aggregate tax preferences when making job offers, one should observe larger distortions in the equilibrium distribution of job offers in sectors or occupations where workers themselves exhibit larger tax elasticities.

We test these three predictions using a matched employer-employee panel of the population in Denmark between 1994 and 2001. This dataset combines administrative records on earnings and taxable income, demographic characteristics, and employment characteristics such as occupation and tenure. There are two sources of tax variation in the data: tax reforms across years, which produce variation in marginal net-of-tax wage rates of 10% or less, and changes in tax rates across tax brackets within a year, which generate variation in net-of-tax wages of up to 35%. We focus primarily on the cross-bracket variation in taxes rates because it is larger and applies to large subgroups of the population, permitting coordinated responses. In particular, we estimate taxable income elasticities by measuring the amount of bunching in earnings at kink points, as in Saez (2009).³

Consistent with the first prediction, the elasticities implied by the amount of bunching at large kinks are significantly larger than those implied by the amount of bunching at smaller kinks. There is substantial, visually evident excess mass in the wage earnings distribution around the cutoff for the top income tax bracket in Denmark, at which the net-of-tax wage rate falls by approximately 30%. There is little excess mass at kinks where the net-of-tax wage falls by 10%, and no excess mass at kinks that generate variation in net-of-tax wages smaller than 10%. Similarly, we find no changes in earnings around the small tax reforms that change net-of-tax wages by less than 10%. The observed elasticities at the largest kinks are 3-5 times those generated by smaller kinks and

³Following the modern public finance literature reviewed in Saez et al. (2009), we proxy for "labor supply" using taxable income. We discuss the implications of measuring taxable income elasticities instead of hours elasticities below.

tax reforms across a broad range of demographic groups, occupations, and years. Using a series of auxiliary tests, we show that the differences in observed elasticities are driven by differences in the size of the tax changes rather than heterogeneity in structural elasticities by income levels or tax rates.

To test the second prediction, we exploit heterogeneity in deductions across workers. In Denmark, 60% of wage earners have zero deductions. These workers reach the top tax bracket when their wage earnings exceeds the top tax cutoff for taxable income, which we term the “statutory” top tax cutoff. Workers with large deductions or non-wage income, however, reach the top tax cutoff at different levels of wage earnings and thus have less common tax incentives. We first demonstrate that firms cater to the tax incentives of the most common workers. In particular, the mode of occupation-level wage earnings distributions has an excess propensity to be located near the statutory top tax cutoff.⁴ Importantly, the wage earnings distribution even for workers who have substantial deductions or non-wage income exhibits excess mass at the *statutory* top tax cutoff. Because these workers do not face any change in marginal tax rates at the statutory cutoff, this finding constitutes direct evidence that firms tailor wage-hours offers to the tax preferences of the majority of workers who have small deductions. We label this firm-driven response to tax incentives “firm bunching.”

Although firm bunching is an important source of behavioral responses to the tax system, some of the bunching at kinks is driven by individual workers searching for jobs that place them near the top tax kink. To isolate and measure such “individual bunching,” we exploit a cap on tax-deductible pension contributions, which is on average DKr 33,000 in the years we study. Approximately 3% of workers make pension contributions up to this amount and therefore cross into the highest income tax bracket when they earn DKr 33,000 more than the statutory top tax cutoff. We find that this pension-driven kink induces excess mass in the distribution of wage earnings at DKr 33,000 above the top tax cutoff. This excess mass appears to be driven solely by individual job search, as there is no excess mass at the pension-driven kink for workers with small deductions. Hence, firms respond only to tax incentives that affect a large group of workers, as the model predicts. Because of firm bunching, workers with common tax preferences (small deductions) have a higher propensity to bunch at the top tax kink than those with uncommon tax preferences (large deductions).

We test the third prediction by estimating the correlation between individual and firm bunching

⁴We focus on wage earnings distributions at the occupation level because most workers’ wages are set through collective bargains at the occupation level in Denmark.

across occupations. We find that firms are more likely to bunch workers at the statutory kink in occupations where workers exhibit more individual bunching in wage earnings at the pension-driven kink. Although this result cannot be interpreted as a causal effect because the variation in the degree of individual bunching is not exogenous, it is consistent with the prediction that firms cater to workers’ tax preferences in equilibrium. Further supporting the importance of firm responses, we find that some of the heterogeneity in elasticities across demographic groups is driven by occupational choice. For instance, reweighting men’s occupations to match those of women’s eliminates 50% of the gap in observed elasticities between men and women.

All of the results above are obtained for wage earners. We analyze self-employed individuals separately. As the self-employed do not face significant adjustment costs or hours constraints, one would expect that none of our three predictions should hold for this subgroup. Indeed, we find that the self-employed exhibit sharp bunching at both small and large kinks, show no evidence of “firm bunching” at the statutory kink, and are equally likely to bunch irrespective of their deductions. These placebo tests support our hypothesis that search costs and hours constraints are the key factors that attenuate micro elasticity estimates for wage earners.

Using the reduced-form empirical evidence described above, we calibrate our model to estimate the structural elasticity ε . As the model has many structural primitives, we develop a method of placing a lower bound on ε without point identifying the remaining parameters. The intuition underlying the bound is that the utility losses agents suffer by deviating from their optimal hours choices are inversely related to ε (Chetty 2009a). A small ε generates significant bunching at *both* small and large kinks because the utility gains from bunching exceed search costs. Hence, for a given level of frictions, ε must exceed a lower bound in order to generate the substantial difference between the elasticities observed at small and large kinks in the data. Implementing this partial identification approach by parametrizing our model, we obtain a lower bound on the structural (or macro) elasticity of $\varepsilon \geq 0.34$.

Micro elasticity estimates understate ε by an order of magnitude in our data. Even at the largest kink where net-of-tax wages fall by 30%, the elasticity implied by the observed amount of bunching (ignoring frictions) is less than 0.02. The observed elasticity is substantially attenuated because the utility loss from ignoring the 30% kink (and optimizing as if there were no increase in tax rates) is less than 2% of consumption if $\varepsilon = 0.34$. Given plausible frictions, the fact that there is any bunching at all implies that the underlying structural elasticity must be significantly larger than 0.02. Micro estimates are attenuated by frictions because they are identified from individuals’

responses to changes in tax rates or kinks *after* obtaining a job near their optimum. In contrast, macro variation in tax rates across countries changes the jobs individuals search for and the jobs offered by firms to begin with, producing larger elasticities.

Our results help explain the longstanding puzzle of why macro studies find much larger elasticities than microeconomic studies. Micro studies have found very small intensive-margin elasticities for all except top income earners (Blundell and MaCurdy 1999, Saez et al. 2009). For instance, Chetty (2009a) reports a mean elasticity estimate of 0.12 based on a meta-analysis of 12 recent studies. In a microeconomic study that uses the same Danish microdata as we do here, Kleven and Schultz (2010) estimate an elasticity of zero by studying tax reforms over a twenty year period. Our elasticity estimates for middle-income wage earners justify calibrating macro models with larger elasticities than these micro estimates. However, we caution that our findings do not provide justification for the very large elasticities (e.g. $\varepsilon > 1$) used in some macro models.

Our explanation for the gap between micro and macro elasticities complements recent work arguing that macro elasticities are larger because they incorporate both extensive and intensive margin responses (e.g. Rogerson and Wallenius 2009). While this insight clearly explains part of the puzzle, much of the difference in labor supply across countries with different tax regimes is driven by hours worked conditional on employment (Davis and Henrekson 2005). That is, macro estimates of *intensive margin* elasticities are much larger than their microeconomic counterparts. Our analysis explains this divergence between intensive margin elasticities.⁵

In addition to the literature on micro vs. macro elasticities, our study builds on and contributes to several other strands of the literature on labor supply. First, previous work has proposed that adjustment costs and hours constraints affect labor supply decisions (e.g. Cogan 1981, Altonji and Paxson 1988, Ham 1982, Dickens and Lundberg 1993). Our contribution is to show how these factors affect estimates of intensive-margin labor supply elasticities. Our findings also support the hypothesis that the effects of government policies may operate through coordinated changes in social norms or institutions rather than individual behavior (e.g. Lindbeck 1995, Alesina, Glaeser and Sacerdote 2005).

Second, our results contribute to the literature on non-linear budget sets (e.g., Hausman 1981, Moffitt 1990, MaCurdy, Green, and Parsch 1990), where the lack of bunching at kinks creates problems in fitting models to the data. As noted by Blundell and MaCurdy (1999), “...for the

⁵See Chetty (2009a) for a more thorough reconciliation of micro and macro elasticities on both the intensive and extensive margins.

vast majority of data sources currently used in the literature, only a trivial number of individuals, if indeed any at all, report [earnings] at interior kink points.” The kinks examined in previous studies are generally much smaller – both in the change in tax rates at the kink and the size of the group of individuals affected – than the largest kinks studied here. Our calibrated model explains behavior at kinks of different sizes by incorporating adjustment costs and hours constraints into a non-linear budget set framework.

Third, our analysis relates to recent work on taxable income as a measure of labor supply. Feldstein (1999) argues that taxable income elasticities are a sufficient statistic for tax policy analysis, but more recent studies argue that it is important to distinguish income shifting from “real” changes in labor supply (Goolsbee 2000, Slemrod and Yitzhaki 2002, Chetty 2009b). We show that the bunching we observe is driven by changes in wage earnings rather than tax avoidance via pension contributions or evasion. However, because our dataset does not contain information on hours of work, we cannot definitively rule out the possibility that some of the responses we observe arise from income shifting. Importantly, distinguishing income shifting from hours of work is not critical for the conclusions we draw here. Although our model focuses on hours choices, its predictions also apply to an environment with adjustment costs and coordination constraints in income shifting.

The paper is organized as follows. In Section II, we set up the model, define micro and macro elasticities formally, and derive the three testable predictions. Section III describes the Danish data and provides institutional background. Section IV presents the empirical results. Section V calibrates the model to bound the macro elasticity using the evidence. Section VI concludes.

II Search Costs and Hours Constraints in a Labor Supply Model

We make two assumptions to tailor our model to the empirical analysis and calibration. First, we only consider intensive-margin labor supply choices and do not model labor force participation. Second, we analyze a static model because our empirical analysis focuses on how search costs and hours constraints interact in equilibrium rather than on the dynamics of adjustment in labor supply. We present some results on responses to tax reforms in a two-period extension of the model in Appendix A, but defer a complete analysis of dynamics with search costs and endogenous hours constraints to future work.

II.A Model Setup

We develop a labor supply model in which firms and workers are price-takers in competitive equilibrium.⁶

Firms. Firms have one-factor linear production technologies. Each firm employs a single worker to produce goods sold at a fixed price p . Let $w(h)$ denote the hourly wage rate paid to workers who work h hours in equilibrium. Firm j posts a job that requires h_j hours of work at the market wage rate $w(h_j)$. We model hours constraints by assuming that a firm cannot change the hours it posts after matching with a worker. This assumption captures the intuition that firms sink capital in a technology that requires a certain amount of labor for production *before* hiring workers.

Firms choose the hours h_j they post to maximize profit:

$$(1) \quad \pi_j = ph_j - w(h_j)h_j$$

Intuitively, firms seek to produce at an hours level where the supply of labor exceeds demand, allowing them to earn profits by paying a wage $w(h_j) < p$. Because firms are free to enter the market at any level of hours h_j , profits are bid to zero, implying that $w(h_j) = w = p$ for all h_j in equilibrium. Let the aggregate distribution of hours required by firms be denoted by a cdf $G(h)$.

Workers. Workers, indexed by i , have quasi-linear utility

$$(2) \quad u_i(c, h) = c - \alpha_i^{-1/\varepsilon} \frac{h^{1+1/\varepsilon}}{1 + 1/\varepsilon}$$

over a numeraire consumption good c and hours of work h . The heterogeneous taste parameter $\alpha_i > 0$, is distributed according to a smooth cdf $F(\alpha_i)$ with full support on a closed interval. This utility specification eliminates income effects and generates a constant wage elasticity of labor supply ε in a frictionless model. We abstract from income effects because the variation in marginal tax rates at kinks that we exploit for identification has little effect on average tax rates and thus generates negligible income effects. We extend the analysis to utility functions that generate non-constant elasticities in Appendix A.

In addition to wage earnings wh_i , each worker also has stochastic non-wage income $y_i \sim F_Y$

⁶We analyze a model with competitive equilibrium so that our only departure from standard neoclassical models is the introduction of frictions. We expect that a model with frictions where wages and hours are determined through bargaining rather than competitive equilibrium would generate the same three testable predictions.

whose realization is unknown at the time she chooses h_i . To characterize tax changes that affect subgroups of the population differently, assume that there are two types of tax systems, indexed by $s = \{NL, L\}$.⁷ Individuals with $s_i = NL$ face a two-bracket non-linear tax system with marginal tax rates of τ_1 and $\tau_2 > \tau_1$. These workers begin to pay the higher tax rate when their incomes $y_i + w_i h_i$ exceed a threshold K . Individuals with $s_i = L$ pay a linear tax rate of τ on all income. With this tax system, individual i has consumption

$$(3) \quad c_i(h_i) = \begin{cases} (1 - \tau_1) \min(y_i + w_i h_i, K) + (1 - \tau_2) \max(y_i + w_i h_i - K, 0) & \text{if } s_i = NL \\ (1 - \tau) (y_i + w_i h_i) & \text{if } s_i = L \end{cases}$$

A fraction ζ of workers face the non-linear tax system NL and the remainder $(1 - \zeta)$ face the linear tax system L . The tax systems workers face are uncorrelated with their tastes: $F(\alpha_i | s_i) = F(\alpha_i)$. Let worker i 's optimal level of hours be denoted by $h_i^* = \arg \max_{h_i} \int c_i(h_i) dF_Y(y_i) - \alpha_i^{-1/\varepsilon} \frac{h_i^{1+1/\varepsilon}}{1+1/\varepsilon}$.

Workers begin their search for a job by drawing an initial offer h_i^0 from the aggregate offer distribution $G(h)$. Each worker can either accept this offer or turn it down and search for another job. A worker who declines her initial offer draw a new offer h_i' from a distribution $G_e(h' | h_i^*)$ centered around her optimal hours choice h_i^* (so that $\mathbb{E}[h' | h_i^*] = h_i^*$). The worker can obtain a more precise draw from G_e by exerting search effort $e \in [0, 1]$: $\text{var}(h_i') = k \cdot (1 - e)$. Exerting e units of search effort has a monetary cost $\phi_i(e)$ for worker i , where ϕ_i is a weakly increasing function.⁸ Workers who choose to decline their initial offer choose search effort e to maximize expected utility net of search costs,

$$\max_e \mathbb{E} u_i(c_i, h_i) - \phi_i(e),$$

and accept a new offer h_i' . This job search process for workers can be viewed as a functional \mathcal{F} that maps an aggregate distribution of hours posted by firms $G(h)$ and wage schedule $w(h)$ to a new distribution $\mathcal{F}(G(h), w(h))$.

Equilibrium. In equilibrium, the labor market must clear at each hours level h at a wage rate $w = p$. Market clearing requires that the distribution of jobs initially posted by firms coincides with the jobs selected by workers after the job search process is complete, i.e. $G(h) = \mathcal{F}(G(h), w)$.⁹ Because of market forces, the hours constraints imposed by firms in equilibrium are endogenous

⁷For example, tax systems often treat single and married individuals differently, in which case the two types in our model would be defined by marital status.

⁸The parametric assumption $\text{var}(h_i') = k \cdot (1 - e)$ is made without loss of generality because $\phi_i(e)$ is unrestricted.

⁹We are unable to obtain general analytical results on existence and uniqueness of $G(h)$, but have found that both of these properties hold for a wide range of parameters using numerical simulations.

to the aggregate distribution of worker preferences. If many workers prefer to work 40 hours per week, many firms choose technologies that allow production with 40 hours of labor per week in equilibrium.

This model, which generates a single wage rate $w = p$, should be viewed as a representation of one sector or occupation in the economy. It is straightforward to generate heterogeneous wage rates by introducing multiple sectors. Suppose there are Q different skill types of workers and Q types of corresponding output goods sold at prices p_1, \dots, p_Q . Workers of type q can only work at firms that produce good q , so there is no interaction across the Q segments of the labor market. Then each sector has an equilibrium wage rate $w_q = p_q$ and an equilibrium hours distribution determined by its workers' preferences according to the model above.

The following sections characterize the properties of the equilibrium hours distribution $G(h)$, focusing on the relationship between tax rates and labor supply. For analytical convenience, we derive the key predictions in a series of special cases. We begin by reviewing the benchmark model without search frictions, the special case in which $\phi_i(e) = 0$.

II.B Special Case 1: Benchmark Frictionless Model

In the frictionless model, the structural preference parameter ε fully determines the effects of taxes on labor supply. This is because workers who face no search costs always choose their unconstrained optimal level of hours h_i^* . For workers with $s_i = L$, who face a linear tax τ , the optimal level of hours is $h_i^* = \alpha_i ((1 - \tau) w)^\varepsilon$. The hours choices of workers who face the non-linear tax system can be characterized analytically if there is no uncertainty about non-wage income ($y_i = 0$). When $y_i = 0$, workers with $s_i = NL$ choose

$$(4) \quad h_i^* = \begin{cases} \alpha_i ((1 - \tau_1) w)^\varepsilon & \text{if } \alpha_i < \underline{\alpha} \\ h_K = \frac{K}{w} & \text{if } \alpha_i \in [\underline{\alpha}, \bar{\alpha}] \\ \alpha_i ((1 - \tau_2) w)^\varepsilon & \text{if } \alpha_i > \bar{\alpha} \end{cases}$$

where $\underline{\alpha} = h_K / ((1 - \tau_1) w)^\varepsilon$ and $\bar{\alpha} = h_K / ((1 - \tau_2) w)^\varepsilon$. Workers with moderate disutilities of labor supply $\alpha_i \in [\underline{\alpha}, \bar{\alpha}]$ bunch at the kink because the net-of-tax wage falls discontinuously at h_K , while workers with low or high disutilities of labor choose hours based on the marginal tax rates in the relevant bracket.¹⁰

¹⁰The logic for why a mass of workers bunch at the kink is captured by the following quote from a Danish construction worker interviewed by the Danish Tax Reform Commission: “By the end of November, some of my colleagues stop working. It does not pay anymore because they have reached the high tax bracket.”

Now consider how variation in the linear tax rate τ affects labor supply. When subject to a higher tax rate, workers of type $s = L$ optimally reduce their work hours by

$$(5) \quad d \log h = \varepsilon \cdot d \log(1 - \tau).$$

This equation shows that the elasticity of hours with respect to the net-of-tax rate $(1 - \tau)$ coincides with the structural parameter ε in the frictionless model. We shall therefore refer to ε as the “structural” elasticity. Workers of type $s = NL$, who are unaffected by τ , do not change hours of work and can be used as a control group in an empirical study. Note that in this one-dimensional model of labor supply, the hours elasticity coincides with the elasticity of taxable wage income (wh) with respect to the net-of-tax-rate: $\varepsilon = \frac{d \log wh}{d \log(1 - \tau)}$. Thus, all the results below apply to both hours and taxable income elasticities.¹¹

The elasticity ε is most commonly estimated using variation in tax rates from tax reforms (Blundell and MaCurdy 1999, Saez et al. 2009). However, ε can also be identified from cross-sectional variation in tax rates using non-linear budget set methods (e.g. Hausman 1981). In particular, the amount of bunching observed at kinks identifies ε (Saez 2009). Let $B = [F(\bar{\alpha}) - F(\underline{\alpha})]$ denote the fraction of type $s_i = NL$ individuals who choose $h_i = h_K$. Under the assumption that there is no uncertainty about non-wage income ($y_i = 0$ for all i) and the approximation that the hours distribution is uniform around the kink, Saez (2009) shows that

$$(6) \quad \varepsilon \simeq \frac{B(\tau_1, \tau_2)/g(h_K)}{K \ln \left(\frac{1 - \tau_1}{1 - \tau_2} \right)} = \frac{b(\tau_1, \tau_2)}{K \ln \left(\frac{1 - \tau_1}{1 - \tau_2} \right)},$$

where $b = B/g(h_K)$ denotes the fraction of workers who bunch at the kink normalized by the density of the hours distribution at the kink. Intuitively, the fraction of individuals who stop working at $h = h_K$ hours because of the drop in the net-of-tax wage at that point is proportional to ε . If workers face uncertainty in unearned income or wage rates, ε can be recovered by measuring the excess mass in a window around the kink and adjusting for the degree of uncertainty, as we demonstrate in section IV.

An important property of equations (5) and (6) is that the observed elasticity coincides with

¹¹Following Feldstein (1999), the modern public finance literature has emphasized that income taxes distort choices beyond hours of work, such as training, effort, and fringe benefits. It is straightforward to incorporate these other margins into the model by assuming that workers have utility over H dimensions of labor supply (h^1, \dots, h^H) and firms post job offers that specify all H dimensions (h^1, \dots, h^H) along with wage rates. In such a model, predictions 1-3 apply with respect to the taxable income elasticity rather than the hours elasticity.

ε irrespective of the magnitude of the change in tax rates or the fraction of workers ζ affected by the tax change.¹² This result underlies microeconomic empirical studies of labor supply that use changes in taxes that affect subgroups of the population to identify ε . We now show that with search costs and hours constraints, observed elasticities vary with the size and scope of tax changes and no longer coincide with ε .

II.C Special Case 2: Search Costs and Worker Responses

In this subsection, we analyze the impact of search costs on behavioral responses to taxation, abstracting from changes in the hours constraints set by firms. To isolate worker responses and obtain analytical results, we specialize the model in three ways. First, we assume that the set of workers affected by the tax change has measure zero. When analyzing bunching at kinks, we assume that the fraction of agents who face the non-linear tax system is $\zeta = 0$; conversely, when analyzing tax reforms, we assume $\zeta = 1$. Under this assumption, the tax change has no impact on the equilibrium offer distribution $G(h)$ and only affects the treated workers' hours through changes in job search. Second, we consider a search cost function that generates a binary search decision: workers either retain their initial hours draw h_0 or pay a search cost ϕ and choose their optimal hours deterministically. That is, we assume $\phi_i(e) = \phi \forall e \in [0, 1]$, so that workers set $e = 1$ if they choose to search. Finally, we assume that there is no uncertainty about non-wage income ($y_i = 0$) as above so that we can measure elasticities from point masses at kinks.

Under these assumptions, a worker searches for a new job if his initial offer $h_i^0 \notin [\underline{h}_i, \bar{h}_i]$, where the thresholds are defined by the equations:

$$(7) \quad u(c_i(h_i^*), h_i^*) - u(c_i(\underline{h}_i), \underline{h}_i) = \phi \text{ with } \underline{h}_i < h_i^*$$

$$(8) \quad u(c_i(h_i^*), h_i^*) - u(c_i(\bar{h}_i), \bar{h}_i) = \phi \text{ with } \bar{h}_i > h_i^*$$

Workers who draw hours that fall within the region $[\underline{h}_i, \bar{h}_i]$ retain their initial offer because the utility gains from working h_i^* hours instead of h_i^0 hours are less than the cost of search ϕ . After the search process is complete, there are two types of workers at each firm j : a point mass whose optimal labor supply $h_i^* = h_j$ is exactly that offered by the firm and a distribution of workers with optimal hours near but not equal to h_j .

Now consider how the mapping from the amount of bunching at kinks to ε in (6) is affected

¹²We use the term ‘‘tax change’’ to refer both to changes in tax rates over time via reforms and changes in marginal tax rates at kinks within a given period.

by search costs. Let $\hat{\varepsilon}(\tau_1, \tau_2) = \frac{B(\tau_1, \tau_2)/g(h_K)}{K \ln\left(\frac{1-\tau_1}{1-\tau_2}\right)}$ denote the elasticity obtained by applying equation (6). We shall refer to $\hat{\varepsilon}$ as the “observed” elasticity from bunching at the kink. To understand the connection between $\hat{\varepsilon}$ and ε , first recall that in the frictionless model (where $\phi = 0$), workers locate at the kink if $\alpha_i \in [\underline{\alpha}, \bar{\alpha}]$. When $\phi > 0$, workers locate at the kink if $\alpha_i \in [\underline{\alpha}, \bar{\alpha}]$ and $h_i^0 \notin [\underline{h}_i, \bar{h}_i]$.¹³ As a result, the observed elasticity $\hat{\varepsilon}$ is smaller than the structural elasticity ε . As the size of the tax change at the kink increases (τ_1 falls or τ_2 rises), the set of workers with $\alpha_i \in [\underline{\alpha}(\tau_1, \tau_2), \bar{\alpha}(\tau_1, \tau_2)]$ who pay the search cost to locate at the kink expands:

$$(9) \quad \frac{\partial[\bar{h}_i - \underline{h}_i]}{\partial\tau_2} < 0 \quad \text{and} \quad \frac{\partial[\bar{h}_i - \underline{h}_i]}{\partial\tau_1} > 0.$$

Because the equilibrium hours distribution $G(h)$ is not affected by τ_1 and τ_2 when $\zeta = 0$, it follows immediately that $\hat{\varepsilon}$ rises with $\tau_2 - \tau_1$. As $\tau_1 \rightarrow -\infty$ and $\tau_2 \rightarrow \infty$, the inaction region $[\underline{h}_i, \bar{h}_i]$ collapses to h_K for agents with $\alpha_i \in [\underline{\alpha}, \bar{\alpha}]$, and $\hat{\varepsilon} \rightarrow \varepsilon$. Intuitively, the effects of large tax changes that affect a measure zero set of workers ($\zeta = 0$) depend purely on workers’ preferences (ε) because they make all the treated workers reoptimize without inducing firm responses.

Larger kinks generate larger observed elasticities because the utility costs of ignoring a kink increase with its size. Figure 1 illustrates this intuition using indifference curves in consumption-labor space for an agent who would optimally set hours at h_K . The thresholds $[\underline{h}_i, \bar{h}_i]$ are where the budget constraint crosses the indifference curve that yields utility ϕ units less than the maximal utility U^* . Now suppose τ_2 increases, moving the upper budget segment from the solid line to the dashed line. Then the upper bound \bar{h}_i decreases, which in turn increases $\hat{\varepsilon}$. This is because the utility loss from supplying hours above the kink rises with τ_2 , as one earns less for this extra effort. These results lead to our first testable prediction about how search costs affect the relationship between taxes and labor supply:

PREDICTION 1: *When workers face search costs, the observed elasticity from bunching rises with the size of the tax change and converges to ε as the size of the tax change grows:*

$$(10) \quad \partial\hat{\varepsilon}/\partial\tau_2 > 0, \quad \partial\hat{\varepsilon}/\partial\tau_1 < 0, \quad \text{and} \quad \lim_{(\tau_2 - \tau_1) \rightarrow \infty} \hat{\varepsilon} = \varepsilon$$

We derive an analogous prediction for observed elasticities from tax reforms in Appendix A.

¹³Workers who draw $h_i^0 \in [\underline{h}_i, \bar{h}_i]$ do not contribute to the point mass at the kink because $G(h)$ is smooth when $\zeta = 0$. Therefore, among type $s = NL$ workers, the set who draw an initial hours offer $h_i^0 = K/w$ has measure zero. $G(h)$ is smooth in this case because the distribution of tastes $F(\alpha)$ is smooth and the set of agents who face a smooth (linear) tax schedule has measure 1.

Tax reforms generate observed elasticities $\hat{\varepsilon} = \frac{d \log h}{d \log(1-\tau)}$ that differ from ε ; as the size of the tax reform grows, $\hat{\varepsilon} \rightarrow \varepsilon$. The intuition for this result is very similar to that for bunching: many workers will not pay the search cost to find a job that requires fewer hours following a tax increase, attenuating $\hat{\varepsilon}$. However, unlike in the case of bunching, observed elasticities from tax reforms need not always be smaller than ε . For example, if workers are close to the edge of their inaction regions prior to the reform – which could in principle occur if there have been a series of small tax increases in the past – then a small tax change could lead to large adjustments, generating $\hat{\varepsilon} > \varepsilon$. Hence, observing that elasticities rise with the size of tax reforms is sufficient, but not necessary, to infer that search costs affect observed elasticities.

Non-Constant Structural Elasticities. If the utility function is not isoelastic, one may observe an elasticity $\hat{\varepsilon}$ that increases with the size of the tax change even without search costs. We can distinguish search costs from variable structural elasticities by comparing the effects of several small tax changes with the effects of a larger change that spans the smaller changes. In Appendix A, we show that with an arbitrary utility $u(c, l)$ and tax rates $\tau_1 < \tau_2 < \tau_3$, the amount of bunching at two smaller kinks is equal to the bunching created at a single larger kink in the frictionless case ($\phi = 0$):

$$B(\tau_1, \tau_3) = B(\tau_1, \tau_2) + B(\tau_2, \tau_3).$$

This is because the amount of bunching increases linearly with the size of the kink without search costs, as shown in (6). In contrast, when $\phi > 0$,

$$B(\tau_1, \tau_3) > B(\tau_1, \tau_2) + B(\tau_2, \tau_3).$$

Intuitively, agents are more likely to pay the fixed search cost ϕ to relocate to the bigger kink, and thus it generates more bunching and a larger observed elasticity than the two smaller kinks together. A similar result applies to tax reforms: the effect of two small tax reforms, each starting from a steady state differs from the effect of one large reform only when $\phi > 0$. We exploit these results to show that the differences in observed elasticities we document in our empirical analysis are driven by search costs rather than changes in the structural elasticity.

Micro vs. Macro Elasticities. Search costs lead to a divergence between the elasticities observed from micro studies of tax reforms or bunching and the elasticities relevant for macroeconomic comparisons. In particular, the structural elasticity ε determines the steady-state effect of variation

in tax policies across economies on aggregate labor supply even with search costs.¹⁴ To see this, consider two economies with different linear tax rates, τ and τ' , for workers with $s_i = L$. To abstract from firm responses to this tax variation, assume that the set of individuals facing the linear tax has measure zero ($\zeta = 1$); we show that the same result holds with firm responses in the next subsection. We define the observed macro elasticity as the effect of this difference in tax rates on hours of work:

$$\widehat{\varepsilon}_{\text{MAC}} = \frac{\mathbb{E} \log h_i(\tau') - \mathbb{E} \log h_i(\tau)}{\log(1 - \tau') - \log(1 - \tau)}$$

For workers who pay the search cost to choose optimal hours, the difference in hours between the two economies is

$$\log h_i^*(\tau') - \log h_i^*(\tau) = \varepsilon \cdot (\log(1 - \tau') - \log(1 - \tau))$$

Workers who retain their original hours draw h_i^0 have average work hours of $\int_{\underline{h}_i}^{\bar{h}_i} h dG(h)$. Under a quadratic approximation to utility, the movement in the inaction region is also determined by ε :

$$\frac{\partial \log \underline{h}_i}{\partial \log(1 - \tau)} = \frac{\partial \log \bar{h}_i}{\partial \log(1 - \tau)} \simeq \varepsilon.$$

Under the approximation that the offer distribution $G(h)$ is uniform between \underline{h}_i and \bar{h}_i ,

$$\mathbb{E} \log h_i(\tau') - \mathbb{E} \log h_i(\tau) \simeq \varepsilon \cdot (\log(1 - \tau') - \log(1 - \tau))$$

It follows that $\widehat{\varepsilon}_{\text{MAC}} \simeq \varepsilon$: the macro elasticity approximately equals the structural elasticity regardless of the search cost ϕ .

The critical difference between micro and macro elasticities is that the former are identified from a worker's decision to switch jobs *ex-post* because of tax incentives, whereas the latter are identified from differences in *ex-ante* job search behavior. Search costs reduce workers' propensity to fine tune their labor supply choices by bunching at kinks or responding to tax reforms because the costs of deviating from optima are second-order. But workers search for jobs with fewer hours to begin with in an economy with higher tax rates. Consequently, a tax reform or a kink that changes the marginal rate from τ_1 to τ'_1 generates a smaller observed elasticity than the *same* "macro" variation in tax rates of τ_1 vs. τ'_1 across economies.

¹⁴Recovering the structural primitives of preferences is also essential for welfare analysis.

II.D Special Case 3: Hours Constraints and Firm Responses

We now show how changes in hours constraints set by firms affect observed responses to tax changes. To highlight firm responses and obtain analytical results, we consider a different special case of the model. First, we assume $\zeta \in (0, 1)$, so that there is a positive measure of workers affected by both tax systems. Second, we assume that at each level of α_i , a fraction δ of workers face no search costs ($\phi_i(e) = \phi_i = 0$) and the remaining workers cannot search at all ($\phi_i(e) = \phi_i = \infty$). Finally, we maintain the assumption above that there is no uncertainty about non-wage income ($y_i = 0$).

In this special case, workers' search decisions are simple: those with $\phi_i = 0$ set $h_i = h_i^*$ and those with $\phi_i = \infty$ set $h_i = h_i^0$, their initial hours draw. As a result, the equilibrium distribution of job offers $G(h)$ coincides with the distribution of optimal hours choices, $G^*(h)$. The reason is that the search process \mathcal{F} maps a distribution of offers to $\mathcal{F}(G) = \delta G^* + (1 - \delta)G$, and hence G^* is the only fixed point of \mathcal{F} . Intuitively, workers with $\phi_i = 0$ always choose their optimal hours, and so the only offer distribution that is a fixed point for them is G^* . As any offer distribution is a fixed point for the $\phi_i = \infty$ group, G^* must be the aggregate hours distribution in equilibrium. This result illustrates that firms cater to workers' preferences when setting hours constraints in equilibrium. Taxes therefore affect labor supply by changing aggregate worker preferences and inducing shifts in hours constraints.

To see how this mechanism affects elasticity estimates, consider the observed elasticity from bunching for the workers who face the non-linear tax ($s_i = NL$). Let $B^*(\tau_1, \tau_2)$ denote the level of bunching that one would observe in the frictionless model ($\delta = 1$) for these workers. With search costs ($\delta < 1$), the observed amount of bunching for workers with $s_i = NL$ is:

$$B = \delta B^* + (1 - \delta)\zeta B^*$$

The two terms in this expression represent two distinct sources of bunching. The first term arises from workers who choose $h_i = h_i^* = h_K$ because they face no search costs. The second term arises from the workers who set $h_i = h_i^0 = h_K$ because they face infinite search costs. Because the aggregate distribution of hours coincides with the optimal aggregate distribution, a fraction ζB^* of the equilibrium job offers have hours of h_K . We label the first component of bunching ($B_I = \delta B^*$) “individual bunching” because it arises from individuals' choices to locate at the kink via job search.¹⁵ We label the second component ($B_F = (1 - \delta)\zeta B^*$) “firm bunching” because it

¹⁵A fraction $(B^*)^2$ of workers with $\phi_i = 0$ and $h_i^* = h_K$ draw the $h_i^0 = h_K$ to begin with and are therefore

arises from workers drawing an initial hours offer that places them at the kink to begin with.

The signature of firm bunching is that it generates bunching even amongst workers who have no incentive to locate at the kink. Consider workers with $s_i = L$, who face a linear tax schedule and experience no change in marginal tax rates at h_K . Because of the interaction of hours constraints with search costs, these workers also bunch at the kink via the firm bunching channel. These workers draw $h_i^0 = h_K$ with probability ζB^* and are forced to retain that offer if $\phi_i = \infty$. The amount of bunching observed for workers with $s_i = L$ is therefore $B_L = (1 - \delta)\zeta B^* = B_F$. This equivalence between B_L and B_F is useful empirically because we cannot measure B_F directly (as we do not observe search behavior), but we can measure B_L since we do observe workers' tax schedules. Intuitively, any bunching among those who do not face a kink must represent firm bunching.

The observed elasticity from bunching for workers with $s_i = NL$ is:

$$\hat{\varepsilon} = \frac{B(\tau_1, \tau_2)/g^*(h_K)}{K \ln\left(\frac{1-\tau_1}{1-\tau_2}\right)} = \delta\varepsilon + (1 - \delta)\zeta\varepsilon < \varepsilon$$

The observed elasticity is smaller than the structural elasticity because search costs prevent some workers who would like to be at the kink from moving there.¹⁶ The observed elasticity rises with the scope of the kink ζ – the fraction of workers in the economy who face the non-linear tax schedule. When more workers face a change in tax incentives at an earnings level of K , firms are compelled to offer more jobs in equilibrium at h_K hours to cater to aggregate preferences. Thus a kink that affects more workers generates more firm bunching (higher B_F) and thereby leads to more total bunching and a larger observed elasticity $\hat{\varepsilon}$.

As the scope of the kink approaches $\zeta = 1$, $B \rightarrow B^*$ and $\hat{\varepsilon} \rightarrow \varepsilon$ in this special case.¹⁷ Conversely, as ζ approaches 0, B_F converges to 0 because firms only cater to aggregate preferences. It follows that the bunching observed at kinks that affect few workers in the economy constitutes a pure

indifferent between retaining h_i^0 and searching for their optimal job. To simplify notation, we classify these workers as “individual bunchers” by assuming that they choose to search for a new job.

¹⁶In this special case, the *total* amount of bunching including all workers (both L and NL) equals the amount of bunching in the frictionless case ($\delta = 0$) because $G(h) = G^*(h)$. However, the composition of those at the kink differs when $\delta > 0$: some of those who bunch face the linear tax. This is why $\hat{\varepsilon} < \varepsilon$ for workers of type NL . In the general model where workers face finite adjustment costs, $G(h) \neq G^*(h)$ and total bunching no longer coincides with that in the frictionless case.

¹⁷The convergence of B to B^* only occurs in this special case. In the general model, $B < B^*$ when $\zeta = 1$ for most parameter values. However, one can sometimes obtain $B > B^*$ when $\zeta = 1$ because of firm bunching. The robust testable prediction is that B rises with ζ .

measure of individual bunching:

$$(11) \quad \lim_{\zeta \rightarrow 0} B = B_I$$

This equivalence between $\lim_{\zeta \rightarrow 0} B$ and B_I is also useful empirically because we cannot directly observe B_I , but can observe $\lim_{\zeta \rightarrow 0} B$ by studying bunching at kinks that apply to few workers.¹⁸ These results lead to our second testable prediction about taxes and labor supply elasticities.

PREDICTION 2: Search costs interact with hours constraints to generate firm bunching. The amount of firm bunching and the observed elasticity rises with the fraction of workers who face the kink:

$$(12) \quad \begin{aligned} B_F &= B_L > 0 \text{ iff } \zeta > 0 \\ \frac{\partial B_F}{\partial \zeta} &> 0 \text{ and } \frac{\partial \hat{\varepsilon}}{\partial \zeta} > 0. \end{aligned}$$

The source of firm bunching is that profit-maximizing firms cater to workers' preferences when setting hours constraints. Therefore, in occupations where workers are more tax elastic, one should observe a higher level of both individual and firm bunching. To see this, consider the Q -sector extension of the model described above. The amount of individual bunching in occupation q is $B_I^q = \delta \zeta B^{q,*}$ and the amount of firm bunching is $B_F^q = (1 - \delta) \zeta B^{q,*}$. As the structural elasticity ε_q increases, the fraction of workers who would optimally locate at the kink ($B^{q,*}$) increases, increasing both B_I^q and B_F^q because δ and ζ are constant.¹⁹ This leads to our third and final prediction.

PREDICTION 3: Firms cater to workers' preferences – the amount of firm bunching and individual bunching are positively correlated across occupations:

$$(13) \quad \text{cov}(B_I^q, B_F^q) > 0$$

In Appendix A, we derive analogs of predictions 2 and 3 for observed elasticities from tax reforms. The analog of “firm bunching” for tax reforms are changes in the hours constraints imposed by firms in response to changes in tax rates. These changes in hours constraints affect hours of work even for workers whose tax incentives are unaffected by the reform, providing an

¹⁸This is why the bunching in special case 2 above (where $\zeta = 0$) is driven purely by individual search behavior rather than firm responses.

¹⁹If workers could switch between sectors, this correlation result would be reinforced because more elastic workers would sort toward sectors with more firm bunching.

empirical measure of firm responses. Tax changes that affect a larger group of workers induce larger changes in hours constraints. As a result, the observed elasticity from a tax reform, $\hat{\varepsilon} = \frac{d \log h}{d \log(1-\tau)}$, increases with the scope of the reform. Firms are more responsive to tax reforms in sectors of the economy where workers are more tax elastic, producing a correlation between observed firm responses and individual responses to tax reforms.

Micro vs. Macro Elasticities. The structural elasticity ε continues to determine the macro elasticity with firm responses. Consider again the two economies with different linear tax rates, τ and τ' , for workers of type $s_i = L$. But now assume that all workers face the linear tax ($\zeta = 0$), so that firms respond to this tax variation. The results above imply that the difference in equilibrium hours across the two economies coincides with the difference in optimal hours. It follows immediately that the difference in average hours of work between the two economies is

$$\mathbb{E} \log h_i(\tau') - \mathbb{E} \log h_i(\tau) = \mathbb{E} \log h_i^*(\tau') - \mathbb{E} \log h_i^*(\tau) = \varepsilon \cdot (\log \tau' - \log \tau)$$

Hence, the observed macro elasticity equals the structural elasticity ($\hat{\varepsilon}_{\text{MAC}} = \varepsilon$) even with firm responses. This result highlights a second reason that the macroeconomic effects of taxes may differ from microeconomic estimates. Variation in tax rates across economies shifts the aggregate distribution of workers' preferences and thereby induces changes in the hours constraints set by firms. In contrast, tax reforms or kinks that affect a small subgroup of workers do not generate substantial changes in hours constraints. Thus, microeconomic studies that focus on tax changes that affect specific groups of the population could significantly underestimate macro elasticities.

We derived the three predictions in special cases because the general model with finite search costs and endogenous hours constraints is analytically intractable. Using numerical simulations, we have verified that the three predictions hold in the general model for parameters that fit the data (see section V). The simulations also confirm that the macro elasticity coincides with ε in the general model. We therefore proceed to test the predictions empirically and determine the extent to which adjustment costs and hours constraints attenuate micro elasticity estimates in practice.

III Institutional Background and Data

The Danish labor market is characterized by a combination of institutional regulation and flexibility, commonly termed “flexicurity.” Virtually all private sector jobs are covered by collective bargaining agreements, negotiated by unions and employer associations. The collective bargains set wages

at the occupation level as a function of seniority, qualifications, degree of responsibility, etc. The contracts are typically negotiated at intervals of 2-4 years. Despite this relatively rigid bargaining structure, rates of job turnover are relatively high and the unemployment rate is relatively low. For example, Anderson and Svarer (2007) report that rates of job creation and job destruction for most sectors and the overall economy in Denmark are comparable to those in the U.S. The unemployment rate in 2000 in Denmark was 5.4%, among the lowest in Europe.

During the period we study (1994-2001), income was taxed using a three-bracket system. Figure 2a shows the tax schedule in 2000 in terms of Danish Kroner (DKr). Note that $\$1 \approx \text{DKr } 6$. The marginal tax rate begins at approximately 45%, referred to as the “bottom tax.”²⁰ At an income of DKr 164,300, a “middle tax” is levied in addition to the bottom tax. The net-of-tax wage rate falls by 11% at the point where the middle bracket begins. Finally, at incomes above DKr 267,600, individuals pay the “top tax” on top of the other taxes, bringing the marginal tax rate to approximately 70%. The net-of-tax wage rate falls by 33% at the point where the top bracket begins. Approximately 25% of wage earners pay the top tax during the period we study. The large jump in marginal tax rates in a central part of the income distribution makes the Danish tax system particularly interesting for our purposes.²¹

Figure 2b plots the movement in the top bracket cutoff across years in real and nominal terms. Danish tax law stipulates that the movement in the top tax bracket from year t to year $t + 1$ is a pre-determined function of wage growth in the economy from year $t - 2$ to year $t - 1$ (two-year lagged wage growth). This mechanical, pre-determined movement of the cutoffs rules out potential concerns that the bracket cutoffs may be endogenously set as a function of labor market contracts. Over the period of study, inflation was between 1.8% and 2.9% per year. Because of the adjustment rule, the top bracket cutoff declines in real terms from 1994-1997 and then increases sharply from 1998-2001.

In addition to the variation in tax rates across brackets, there were also some small tax reforms during the period we study. For example, before 1996, there were two separate middle taxes that were consolidated into a single middle tax in subsequent years. Starting in 1999, capital losses could not be deducted from the middle tax base and contributions to certain types of pensions could no longer be deducted from the top tax base. Finally, the middle and top tax bracket cutoffs

²⁰Individuals with very low incomes (3% of wage earners) are exempt from this bottom tax.

²¹Denmark also has a complex transfer system that affects incentives for low incomes (Kleven and Kreiner 2006). We do not model the transfer system here because transfer programs affect very few individuals’ marginal incentives around the middle and top tax cutoffs that are the focus of our empirical analysis.

change in real terms across years. These tax reforms generate changes in net-of-tax rates between -10% to +10% for certain subgroups of the population, yielding several tax changes of small size and scope.

We study behavior at the individual level because spouses file individual tax returns, but there are some joint aspects of the tax system, which we account for in our analysis. There are two tax bases relevant for our analysis: one for the top tax and one for the middle taxes. We use the term “taxable income” to refer to the tax base relevant to a particular tax; for instance, when studying bunching around the top tax cutoff, we use “taxable income” to refer to the top tax base.²² Wage earnings, self-employment income, transfer payments, and gifts are all subject to both the middle and top income taxes. Most pension contributions are tax deductible and the marginal dollar of capital income is not subject to the top tax for most individuals. These features of the tax code create an incentive to shift earnings from labor income to capital income and pensions. See Ministry of Taxation (2002) for a more comprehensive description of the Danish tax system.

Data. We merge several administrative registers provided by Statistics Denmark. The primary dataset is the tax register from 1994-2001, which contains panel data on wage earnings, self-employment income, pensions, capital income and deductions, spouse ID, and several other characteristics. The tax register contains records for more than 99.9% of individuals between the ages of 15-70 in the population. We merge the tax data with the Danish Integrated Database for Labor Market Research (IDA), which includes data on education, firm ID, occupation, labor market experience, and number of children for every person in Denmark. Additional details on the dataset and variable definitions are given in Appendix B.

Starting from the population dataset, we restrict attention to individuals who (1) are between the ages of 15 and 70 and (2) are wage earners, excluding the self-employed and pensioners.²³ These exclusions leave us with an analysis sample of 17.9 million observations of wage earners. Much of our analysis focuses on the subset of 6.8 million observations for wage earners that fall within DKr 50,000 of the top tax cutoff. We also study the 1.8 million observations of self-employed individuals separately.

Table 1 presents summary statistics for the population of 15-70 year olds as a whole, all wage earners, the subset of wage earners within DKr 50,000 of the top tax cutoff, and self-employed

²²The Danish tax system includes a technical concept of “Taxable Income.” Our use of the term “taxable income” does not refer to that technical concept.

²³The endogenous sample selection induced by dropping the self-employed does not spuriously generate bunching. There is significant bunching in the wage earnings distribution even in the full population: $b = 0.73$ in the full population vs. $b = 0.71$ for the subgroup of wage earners reported in Figure 7 below.

individuals. The mean individual personal (non-capital) income in the population is DKr 180,213 (\$30,000) for the population and DKr 227,359 (\$38,000) for wage earners. Mean net capital income is negative because mortgage interest payments exceed capital income for most individuals. We define “net deductions” as deductions minus non-wage income (accounting for spousal deductions), or equivalently, wage earnings minus taxable income. Most wage earners have small net deductions (60% have deductions less than DKr 7,500 in magnitude), a fact that proves useful for our empirical analysis. The mean level of net deductions is negative because some individuals have substantial non-wage income.

We construct a tax simulator that calculates tax liabilities and marginal tax rates using these data. The tax simulator predicts actual tax liabilities within DKr 5 (\$1) for 95% of the individuals in the population. Over the period we consider, top marginal tax rates were reduced slightly, and thus the simulated net-of-tax rate (holding fixed base-year characteristics) rises by 2.25% on average across two-year intervals.

IV Empirical Analysis

We begin by analyzing bunching at the top bracket cutoff, where net-of-tax wages fall by approximately 30%. Figure 3 plots the empirical distribution of taxable income for all wage earners in Denmark from 1994-2001. To construct this histogram, we first calculate the difference between the actual taxable income and the taxable income needed to reach the top tax bracket for each observation. We then group individuals into DKr 1,000 bins (-500 to 500, 500 to 1500, etc.) on this recentered taxable income variable. Finally, we plot the bin counts around the top bracket cutoff, demarcated by the red vertical line at zero.

Figure 3 shows that there is a spike around the top bracket cutoff in the otherwise smooth and monotonically declining income distribution. As shown in equation (6), the observed elasticity $\hat{\epsilon}$ implied by this bunching is proportional to $b(\tau_1, \tau_2)$, the excess mass relative to the density around the kink K . A complication in measuring b empirically is that noise in non-wage income y_i leads to diffuse excess mass around K rather than a point mass at K . To measure b in the presence of such noise, we must estimate a counterfactual density – what the distribution would look like if there were no change in the tax rate at K . To do so, we first fit a polynomial to the counts plotted

in Figure 3, *excluding* the data near the kink, by estimating a regression of the following form:

$$(14) \quad C_j = \sum_{i=0}^q \beta_i^0 \cdot (Z_j)^i + \sum_{i=-R}^R \gamma_i^0 \cdot \mathbf{1}[Z_j = i] + \varepsilon_j^0$$

where C_j is the number of individuals in income bin j , Z_j is income relative to the kink in 1,000 Kroner intervals ($Z_j = \{-50, -49, \dots, 50\}$), q is the order of the polynomial, and R denotes the width of the excluded region around the kink (measured in DKr 1,000). We define an initial estimate of the counterfactual distribution as the predicted values from this regression omitting the contribution of the dummies around the kink: $\widehat{C}_j^0 = \sum_{i=0}^q \widehat{\beta}_i^0 \cdot (Z_j)^i$. The excess number of individuals who locate near the kink relative to this counterfactual density is $\widehat{b}_n^0 = \sum_{j=-R}^R C_j - \widehat{C}_j^0 = \sum_{i=-R}^R \widehat{\gamma}_i^0$. This simple calculation overestimates b_n because it does not account for the fact that the additional mass at the kink comes from points to the right of the kink; that is, it does not satisfy the constraint that the area under the counterfactual must equal the area under the empirical distribution. To account for this problem, we shift the counterfactual distribution to the right of the kink upward until it satisfies the integration constraint. In particular, we define the counterfactual distribution $\widehat{C}_j = \widehat{\beta}_i \cdot (Z_j)^i$ as the fitted values from the regression

$$(15) \quad C_j \cdot (1 + \mathbf{1}[j > R] \frac{\widehat{b}_n}{\sum_{j=R+1}^{\infty} C_j}) = \sum_{i=0}^q \beta_i \cdot (Z_j)^i + \sum_{i=-R}^R \gamma_i \cdot \mathbf{1}[Z_j = i] + \varepsilon_j$$

where $\widehat{b}_n = \sum_{j=-R}^R C_j - \widehat{C}_j = \sum_{i=-R}^R \widehat{\gamma}_i$ is the excess mass implied by this counterfactual.²⁴ Finally, we define our empirical estimate of b as the excess mass around the kink relative to the average density of the counterfactual earnings distribution between $-R$ and R :

$$(16) \quad \widehat{b} = \frac{\widehat{b}_n}{\sum_{j=-R}^R \widehat{C}_j / (2R + 1)}$$

The solid curve in Figure 3 shows the counterfactual density $\{\widehat{C}_j\}$ predicted using this procedure with a seventh-degree polynomial ($q = 7$) and a window of DKr 15,000 centered around the kink ($R = 7$). The shaded region shows the estimated excess mass around the kink. With these parameters, we estimate $b = 0.81$ – the excess mass around the kink is 81% of the average height

²⁴Because \widehat{b}_n is a function of $\widehat{\beta}_i$, the dependent variable in this regression depends upon the estimates of $\widehat{\beta}_i$. We therefore estimate (15) by iteration, recomputing \widehat{b}_n using the estimated $\widehat{\beta}_i$ until we reach a fixed point. The bootstrapped standard errors that we report below adjust for this iterative estimation procedure.

of the counterfactual distribution within DKr 7,500 of the kink. We choose $q = R = 7$ based on the numerical simulations of the calibrated model described in section V below. In particular, the estimated \hat{b} using (16) is within 15% of the true value in the frictionless model for $b \in [0, 5]$ when $F_Y(y_i)$ is calibrated to match the variance of non-wage income in the data. The qualitative results we report below are not sensitive to changes in q and R or the way in which we correct the counterfactual to satisfy the integration constraint.²⁵ The reason is that the differences we document in observed elasticities are much larger than the changes induced by varying the specification of the counterfactual.

We calculate a standard error for \hat{b} using a parametric bootstrap procedure. We draw from the estimated vector of errors ε_j in (15) with replacement to generate a new set of counts and apply the technique above to calculate a new estimate \hat{b}^k . We define the standard error of \hat{b} as the standard deviation of the distribution of \hat{b}^k s. Since we observe the exact population distribution of taxable income, this standard error reflects error due to misspecification of the polynomial for the counterfactual income distribution rather than sampling error. In Figure 3, the standard error associated with our estimate of b is 0.05. The null hypothesis that there is no excess mass at the kink relative to the counterfactual distribution is rejected with a t-statistic of 17.6, implying $p < 1 \times 10^{-9}$.

There is substantial heterogeneity across groups in the amount of bunching. Panel A in Figure 4 shows that excess mass at the kink is much larger for married women ($b = 1.79$) than for single men ($b = 0.25$), consistent with existing evidence that married women exhibit the highest labor supply elasticities. Panel B shows that there is also substantial heterogeneity across occupations: teachers exhibit substantial bunching around the kink ($b = 3.54$), whereas the military does not ($b = -0.12$, statistically insignificant).²⁶ We return to explore the sources of this heterogeneity in Section IV.C below.

The identification assumption underlying causal inference about the effect of taxes on earnings in the preceding analysis is that the income distribution would be smooth if there were no jump in tax rates at the location of the top bracket cutoff. This identification assumption can be relaxed by exploiting the movement in the top bracket cutoff across years. Figure 5 displays the distribution of taxable income in each year from 1994-2001 for all wage earners and for married women. The

²⁵For example, we obtain similar results by shifting the counterfactual rightward instead of upward to satisfy the integration constraint. Any of these counterfactual adjustments have small effects on the excess mass calculations. Even our initial unadjusted estimate \hat{b}_n^0 differs from the adjusted estimate \hat{b}_n by less than 10%.

²⁶Approximatley 50% of wage earners in Denmark work in the public sector. We find slightly more bunching for those employed in the private sector ($b = 0.67$) than those in the public sector ($b = 0.5$).

excess mass for both groups follows the movement in the top bracket cutoff very closely. In Figure 6, we investigate whether the excess mass tracks tax changes, inflation, or average wage growth over time. Figure 6a considers the period from 1994 to 1997, during which the top tax cutoff declines in real terms. Noting that the excess mass is located at the top tax cutoff in 1994, the figure shows three possibilities for its location in 1997: the 1997 top tax cutoff, the 1994 cutoff adjusted for inflation, and the 1994 cutoff adjusted for average wage growth in the economy. In both the full population of wage earners and the subgroup of married women, the excess mass at the 1994 kink clearly moves to the 1997 kink rather than following inflation or average wage growth. Figure 6b replicates Figure 6a for the period from 1997 to 2001, during which the top tax cutoff rises in real terms. In this figure, the dashed vertical lines show the 1997 top tax cutoff adjusted for inflation and wage growth. Again, the excess masses clearly follow the movement in the top tax cutoff rather than inflation or average wage growth. We conclude that earnings dynamics around the top tax bracket depart from prevailing inflation patterns and instead are aligned with changes in the tax system. We show that firm responses explain why the excess mass tracks the movement of the kink so closely despite frictions in Figure 17 below.

Shifting vs. Real Responses. Individuals can obtain taxable income near the top bracket cutoff through two margins: changes in labor supply (e.g. hours worked) or “income shifting” responses such as changes from taxed to untaxed forms of compensation. Our three theoretical predictions about how frictions affect observed taxable income elasticities hold regardless of what margins underlie changes in taxable income. Intuitively, if firms face technological constraints that limit the benefit packages workers can choose from, tax changes of larger size and scope will continue to produce larger taxable income elasticities. Nevertheless, it is useful to distinguish between these two behavioral responses because income shifting and “real” changes in labor supply have different normative implications (Slemrod and Yitzhaki 2002, Chetty 2009b).

There are two channels through which individuals can change their reported taxable income without changing labor supply: evasion and avoidance. Kleven et al. (2009a) conduct an audit study of Danish tax records and find that there is virtually no tax evasion in wage earnings because of third-party reporting by firms. The lower series in Figure 7 plots the distribution of wage earnings, defined using the same line on the tax form as the wage earnings variable audited by Kleven et al. It shows that there is substantial bunching ($b = 0.68$) even in this narrowest, double-reported measure of compensation. We therefore conclude that the bunching we observe is not driven by evasion.

The second and more important income shifting channel is legal tax avoidance. The simplest method of reducing current tax liabilities is to contribute to tax-deductible pension accounts. We investigate the extent of such shifting by adding employer and employee pension contributions back to taxable income. The upper series in Figure 7 plots the distribution of this broader measure of compensation relative to the statutory top tax bracket cutoff that would apply to individuals with *zero* pension contributions. There is still considerable bunching at the top kink, rejecting the hypothesis that all of the bunching observed in taxable income is driven by shifts to pensions. The excess mass around the top bracket is now smaller than in Figure 3: $b = 0.48$ for pensions plus taxable income, compared with $b = 0.81$ for taxable income. This is not surprising because the vertical line at zero no longer represents the point at which tax rates jump for individuals who make pension contributions. For such individuals, there is no reason to locate at this point even if they are only changing labor supply. To correct for this mechanical attenuation effect, we rescale the estimated excess mass by the fraction of individuals who make pension contributions of more than Dkr 7,500 (which would pull them out of the window we use to compute excess mass). This fraction is approximately 30%, implying an adjusted $b = 0.48/(1 - 0.3) = 0.69$. We conclude that pension shifting is responsible for less than 15% of the bunching in taxable income observed at the top tax cutoff. The relatively small amount of pension shifting is likely driven by the generosity of Denmark’s social security programs. An analogous exercise shows that shifting into capital income, which is untaxed in the top tax base, is responsible for virtually none of the bunching at the top kink.

Although the behavioral responses at the top tax cutoff do not appear to be driven by any observable method of income shifting, we cannot rule out the possibility that individuals shift their compensation to unobservable nontaxable compensation to avoid paying the top income tax. For example, we cannot detect substitution of compensation from wage earnings into office amenities when individuals cross into the top tax bracket. We also cannot rule out intertemporal shifting of wage earnings to avoid paying the top tax. The only way to definitively rule out such responses is to examine changes in hours worked directly. Unfortunately, our dataset does not contain information on hours of work. Nevertheless, we believe that most of the observed bunching in taxable income reflects “real” distortions in behavior that have efficiency costs. Few salaried workers at the 75th percentile of the income distribution have the ability to shift income into other forms of compensation or across time (Slemrod 1995, Goolsbee 2000). Moreover, even if compensation is distorted toward office amenities instead of wages, the marginal efficiency cost of

such distortions equals the marginal efficiency cost of changes in hours of work (Feldstein 1999).

IV.A Prediction 1: Size of Tax Changes

We now test the first prediction by comparing the amount of bunching at the top tax kink with bunching at smaller kinks and observed elasticities from small tax reforms. Figure 8 shows the distributions of taxable income around the middle tax cutoff, where the net-of-tax rate falls by between 7% and 10%.²⁷ Panel A of Figure 8 shows that there is virtually no bunching at the middle tax cutoff ($b = 0.06$) in taxable income for the full population of wage earners. Moreover, the estimated excess mass at the middle tax converges to zero as the degree of the polynomial is increased, whereas the estimated excess mass at the top kink is not sensitive to the degree of the polynomial. Because the tax base for the middle tax differs slightly from that of the top tax, in Panel B we plot the distribution of wage earnings around the middle tax cutoff. There is no evidence of bunching in wage earnings, unlike at the top tax cutoff. Panels C and D show that the amount of bunching remains small and statistically insignificant even for the subsample of married women, who exhibit substantial bunching at the top kink as shown in Figure 4a.

Note that smaller kinks should generate less bunching even in the frictionless model, simply because the change in incentives is smaller. We therefore compare the excess mass at these smaller kinks with the amount of excess mass that would be generated if the elasticity were the same as that implied by the excess mass at the large top tax kink. In all cases, the amount of bunching observed in the empirical distribution at the middle kink is significantly less than what would be predicted by the frictionless model. For example, the frictionless model predicts $b = 0.16$ at the middle kink for all wage earners (Panel A). The null hypothesis that the predicted excess mass equals the actual excess mass at the middle kink can be rejected with $p < 0.01$.

Next, we estimate observed elasticities using changes in marginal rates by legislated reforms. As described in Section III, there were a number of small tax reforms in Denmark between 1994 and 2001 that created changes in net-of-tax rates of between -10% and +10%. These reforms generate differential changes in net-of-tax rates across income groups, motivating a difference-in-difference research design that is the benchmark in the taxable income literature (Saez et al. 2009). Let $\Delta \log y_{i,t} = \log y_{i,t} - \log y_{i,t-2}$ denote the log change in wage earnings from period $t - 2$ to t and $\Delta \log(1 - MTR_{i,t})$ the log change in net-of-tax rates over the same period. Following Gruber and

²⁷In 1994 and 1995, the tax system includes an additional “upper middle tax.” Figure 8 only considers the lower middle tax in these years, but there is no bunching at the upper middle tax cutoff either.

Saez (2002), we estimate the following regression specification using two-stage-least-squares:

$$(17) \quad \Delta \log y_{i,t} = \alpha + \beta \Delta \log(1 - MTR_{i,t}) + f(y_{i,t-2}) + \gamma X_{i,t-2} + \varepsilon_{i,t},$$

instrumenting for $\Delta \log(1 - MTR_{i,t})$ with $\Delta \log(1 - MTR_{i,t}^{\text{sim}})$, the simulated change in net-of-tax rates holding the individual’s income and other characteristics fixed at their year $t - 2$ levels. The function $f(y_{i,t-2})$ is a 10 piece linear spline in base year wage earnings and the vector $X_{i,t-2}$ is a set of base year controls that we vary across specifications. First-stage regressions of $\Delta \log(1 - MTR_{i,t})$ on $\Delta \log(1 - MTR_{i,t}^{\text{sim}})$ have coefficients of approximately 0.6 with F-statistics exceeding 400. The average size of the tax changes used to identify β in (17) is $\mathbb{E}|\Delta \log(1 - MTR_{i,t}^{\text{sim}})| = 4\%$.

Table 2 reports TSLS estimates from several variants of (17). In column 1, we estimate (17) on the full population of wage earners with the following controls: the 10-piece wage earnings spline, a 10-piece spline in total personal income and age and year fixed effects. The estimated elasticity $\hat{\varepsilon}$ is very close to 0, and the upper bound of the 95% CI is $\hat{\varepsilon} = 0.001$. Column 2 adds a 10 piece capital income spline, gender and marital status dummies, and occupation and region fixed effects as controls. The estimated elasticity remains very close to zero, showing that the estimates are robust to the set of covariates used to predict income growth. Column 3 considers the subgroup of married women using the baseline specification in column 1. The observed elasticity in response to small tax changes remains near 0 for married women despite the fact that they exhibit substantial bunching at the large top tax kink, as shown in Figure 4. In column 4, we further restrict the sample to married women who are professionals and have above-median (more than 19 years) labor market experience. This subgroup also does not react significantly to small tax reforms, yet it exhibits substantial bunching at the top kink ($b = 4.50$, with an implied observed elasticity of 0.06). These results match those of Kleven and Schultz (2010), who estimate a similar set of specifications and find near-zero observed elasticities using tax reforms spanning a larger set of years in Denmark.

Figure 9 compiles the evidence on observed elasticities and the size of tax changes by plotting observed elasticities ($\hat{\varepsilon}$) vs. the change in the net-of-tax-rate $\Delta \log(1 - \tau)$ used for identification. We convert the excess mass (b) estimated at the middle and top tax kinks analyzed above into observed elasticities using equation (6). The change in the net-of-tax rate at the middle tax cutoff ranged from 8.9%-9.9% in 1994-1996 vs. 11.4%-11.7% from 1997-2001. We therefore estimate separate excess masses and observed elasticities at the middle tax cutoff for these two sets of years. Similarly, we estimate two observed elasticities for the top tax cutoff: one pooling 1994, 1997,

and 1998 (changes in NTR from 28.1%-30.7%) and another pooling 1995, 1996, and 1999-2001 (31.9%-35.4%). The figure also shows the elasticity estimate using small tax reforms (which have an average size of $\Delta \log(1 - \tau) = 4\%$), using the estimate from Column 1 of Table 2. Recall that the benchmark frictionless model predicts that $\hat{\varepsilon}$ does not vary with $\Delta \log(1 - \tau)$. To test this hypothesis, we fit a linear regression to the five observed elasticity estimates. The hypothesis that the slope of the regression line equals zero is rejected with a t-statistic of 3.6 ($p < 0.05$), supporting prediction 1. The same upward sloping relationship holds across various age groups, years, regions, and occupations.

Although the observed elasticities rise with the size of the tax change, the elasticity implied by the frictionless model remains very small even at the largest kink. For all wage earners, the observed elasticity from bunching at the 30% kink is $\hat{\varepsilon} \simeq 0.01$, while for married women it is $\hat{\varepsilon} \simeq 0.02$. In section V, we calibrate the model and show that the structural (or macro) elasticity that matches the estimates in Figure 9 is $\varepsilon \geq 0.34$, an order of magnitude larger than the observed elasticity at the top kink.

Search Costs vs. Non-Constant Elasticities. If $\varepsilon(\tau, z)$ varies with τ or z , the upward sloping relationship in Figure 9 could potentially be due to variation in ε rather than adjustment costs. In our application, the middle kinks are at incomes of DKr 130,000-177,900, while the top kinks are at incomes of DKr 234,900-276,900. If higher income individuals are more elastic, one would observe the pattern in Figure 9 even without frictions. We distinguish this explanation of our findings from frictions using three approaches.

First, we build upon the results in Section II.C and analyze the effects of small tax changes on the subset of high income individuals who face the top tax to begin with. If the local structural elasticity $\varepsilon(\tau, z)$ around the top kink is larger than $\varepsilon(\tau, z)$ around the middle kinks, then a small change in the top tax rate should generate large behavioral responses. But if the difference in observed elasticities is due to frictions, this small tax change should not generate a significant response. Columns 5 and 6 of Table 2 show estimates of the observed elasticity for all wage earners and married women using the baseline specification in column 1, restricting attention to those with wage earnings exceeding DKr 200,000. The estimated elasticities remain very close to zero. Even conditional on initial income and marginal tax rates, small tax changes induce little or no behavioral response, while larger tax changes induce much larger responses.

Second, we examine how the degree of bunching changes as the middle and top tax cutoffs move across years. In the latter years of our sample, the middle tax cutoff is higher in the income

distribution, but the amount of bunching remains near zero (not shown). In contrast, bunching at the top kink remains substantial across all the years as the bracket cutoff moves up, as shown in Figure 5. These results suggest that the heterogeneity in elasticities across income levels is not significant in the range we study.

As a third test of whether preference heterogeneity drives the differential bunching at the middle and top kinks, we focus on a subset of individuals whose incomes place them within DKr 50,000 of the top kink in year t and within DKr 50,000 of the middle kink in year $t + 2$. By studying these “switchers,” we can effectively remove individual fixed effects when comparing responses to the middle and top kinks. Figure 10 displays the distributions of taxable income for the group of switchers at each kink. When near the top kink, these individuals exhibit substantial bunching ($b = 0.54$). However, just two years later, the same individuals show no excess propensity to bunch at the middle kink ($b = 0.06$) despite having earnings near that kink. The opposite pattern is observed for those moving from the middle to the top kink (not shown). We conclude that variation in $\varepsilon(\tau, z)$ is unlikely to explain the upward-sloping pattern in Figure 9.

Perceptions of the Middle vs. Top Cutoffs. In Chetty et al. (2010), we report the results of an internet survey that asked individuals to report their best guess of the top tax and middle tax cutoffs in the current year. The survey, conducted from March-April 2009, was administered to members of a union representing public and financial sector employees (FTF-A). Our questions were attached to the end of a longer survey on perceptions of the union. 3,299 individuals (11% of the union members) responded to our questions.

Figure 11 displays the distribution of respondents’ perceptions of the middle and top tax cutoffs. Each point depicts the fraction of responses in a DKr 30,000 bin centered around the true cutoffs, shown by the vertical lines. Knowledge of the top tax cutoff is better than the middle tax cutoff. Approximately 41% of those surveyed know the top tax cutoff to within DKr 15,000 of the true level; in contrast, only 27% know the middle tax cutoff with this level of accuracy. The median absolute error for the top tax cutoff is DKr 21,200, compared with DKr 31,000 for the middle tax cutoff. The same qualitative pattern is exhibited across all education levels and occupations in the sample. These survey responses must be viewed as anecdotal evidence because the survey was administered only to members of FTF-A and because the response rate for the questions we added is only 11%. Nevertheless, this evidence is consistent with our finding that observed elasticities are larger at the top kink than the middle kink, as well as recent evidence that the information and salience affect behavioral responses to income taxation (e.g. Chetty and Saez 2009).

IV.B Prediction 2: Firm Bunching and Scope of Tax Changes

To test the second prediction, we begin by identifying a source of variation in the scope of kinks – the fraction of workers in the economy who face a given kink in the tax system. Recall that taxable income is the sum of wage earnings and non-wage income minus deductions. Deductions consist primarily of pension contributions. Non-wage income includes items such as alimony receipts, stipends, and unemployment benefits. Because of heterogeneity in non-wage income and deductions, the wage earnings required to reach the middle and top brackets vary across individuals.

Figure 12a shows a histogram of net deductions (deductions minus non-wage income) in the top tax base. Approximately 60% of wage earners have net deductions less than DKr 7,500 in magnitude, as shown in Table 1. This is because most individuals in Denmark make no pension contributions and earn only wage income. Thus, most individuals cross into the top tax bracket when their wage earnings exceed the top tax cutoff that applies to taxable income, which we term the “statutory” top tax cutoff. The distribution of deductions for the remaining 40% of individuals is diffuse, with one exception, highlighted in Figure 12b. This figure shows the distribution of deductions conditional on having deductions greater than DKr 20,000. There is a spike in the distribution of deductions at DKr 33,000, demarcated by the solid vertical line, which is driven by a cap on tax-deductible pension contributions at DKr 33,000.²⁸ These individuals, who constitute 2.7% of wage earners, reach the top tax bracket only when their wage earnings exceed the statutory top tax cutoff by DKr 33,000.

In this setting, the second prediction of our model consists of three parts: we should observe (1) significant firm bunching at the statutory top tax kink that applies to 60% of workers, (2) little firm bunching at the “pension kink” that applies to 2.7% of workers, and (3) more bunching for individuals with small deductions, as they have more common tax preferences. To test these hypotheses, we study wage earnings distributions at the occupation level because most wages are set through collective bargains at the occupation level in Denmark.

Firm bunching is easiest to see through case studies of occupations. Consider school teachers, who constitute approximately 3% of wage earners in Denmark and form one of the largest unions. Figure 13a plots the distribution of wage earnings around the top tax bracket for teachers. There is very sharp bunching around the statutory top tax cutoff, consistent with the sharp bunching in taxable income shown in Figure 4b.²⁹ Intuitively, the rate of return to negotiating for higher wages

²⁸The pension contribution cap increases slightly over time. In Figure 12b, the distribution of deductions has been recentered so that the cap falls at 33,000 in each year.

²⁹The smaller peak above the kink is driven by teachers in Copenhagen, who receive a cost-of-living adjustment

falls discontinuously for the vast majority of teachers at the top tax bracket cutoff. It is therefore sensible that teachers start bargaining on other dimensions, such as lighter teaching loads or more vacations, rather than continue to push for wage increases beyond this point.

Figure 13b plots the distribution of wage earnings (salaries) around the statutory top tax cutoff for teachers with net deductions *greater* than DKr 20,000. The individuals in this figure do not begin to pay the top tax on wage earnings until at least DKr 20,000 beyond the statutory top tax cutoff, and therefore experience no change in net-of-tax wages at the vertical line at zero. Yet the wage earnings distribution for these workers is extremely similar to the distribution for teachers as a whole, and exhibits sharp bunching at the *statutory* top tax cutoff. This is the signature of firm bunching: even individuals who are unaffected by a kink bunch there. Intuitively, school districts are forced to offer a limited number of wage-hours packages in order to coordinate class schedules. Because of such technological constraints, teachers' contracts cater to the *aggregate* tax incentives reflected in the population. The fact that many teachers would prefer a salary that places them near the statutory kink distorts the earnings of the minority of teachers whose tax incentives differ.

There are similar patterns of firm bunching in many other occupations. We generalize from such case studies by analyzing the modes of the earnings distribution in each occupation, defined using the International Labour Organization's 4 digit International Standard Classification of Occupations (ISCO) codes. We define the mode in each occupation-year cell as the DKr 5,000 wage earnings bin that has the largest number of workers. Figure 14 shows a histogram of these modes relative to the top tax bracket cutoff, excluding small occupation-years that have less than 7,000 workers (25% of the sample). The density of modes drops sharply at the top tax threshold. There are 20 modes within DKr 2000 of the top tax cutoff, but only 6 in the adjacent bin from DKr 2,000 to DKr 6,000 above the kink. This drop in the frequency of modes across these two bins is larger than any other drop across two contiguous bins in the figure. Moreover, as the top tax cutoff rises over years, the distribution of modes shifts along with the cutoff (not shown). Hence, aggregate tax incentives – which are determined largely by the preferences of workers who face the statutory cutoff – shape the distribution of jobs offered by firms as the model predicts.

Having established the prevalence of firm bunching at the most common kink, we test whether kinks that affect fewer workers generate less firm bunching. To do so, we exploit the “pension kink” described above. Figure 15a plots the distribution of wage earnings relative to the pension

of DKr 15,000 over the base teacher's salary. The setting of salaries to place teachers outside Copenhagen – who account for 75% of all teachers – at the top kink supports the view that institutional constraints are endogenously set based on the preferences of the largest groups in the population.

kink (shown by the vertical line at 0) for individuals who have deductions greater than DKr 20,000. There is significant bunching in wage earnings at the pension kink ($b = 0.70$).³⁰ To investigate whether this bunching is driven by firm offers or individual job search, Figure 15b replicates 15a for workers with deductions between DKr 7,500 and DKr 25,000. Note that these workers' tax incentives change at neither the statutory kink nor the pension kink. These workers exhibit no excess propensity to locate near the pension kink ($b = -0.04$), implying that there is little firm bunching at the pension kink. In contrast, Figure 15c shows that the *same* workers exhibit substantial bunching around the statutory kink ($b = 0.58$), confirming that there is significant firm bunching at the statutory kink. Together, these figures offer two lessons. First, the bunching at the pension kink is driven by individual job search – i.e., finding a job that pays DKr 33,000 above the top kink – rather than distortions in the distribution of offers.³¹ Second, firm bunching is significant only at kinks that affect large groups of workers, consistent with the model's prediction that firms cater to aggregate worker preferences.

We now turn to the third part of prediction 2: do workers with small deductions bunch more than those with large deductions? The econometric challenge in testing this prediction is that deductions themselves are endogenous. In particular, workers with large deductions may have chosen their deductions in order to reach the top tax kink.³² We address this endogeneity problem using a grouping instrument. We compute the fraction of workers with deductions less than DKr 7,500 in magnitude for cells of the population defined by marital status, gender, year, and age (in decades). We then divide workers into ten equal-width bins based on the fraction of workers with small deductions in their group and estimate the degree of bunching at the top kink (b) for workers in each of these ten bins.³³ Figure 16 plots the estimated b vs. the fraction of workers with small deductions in the ten groups. The groups with small deductions exhibit much greater bunching: the slope of the fitted line in Figure 16 is statistically significant with $p < 0.01$. This

³⁰We condition on having deductions greater than DKr 20,000 to isolate the relevant part of the population in order to detect bunching at the pension kink. To allay the concern that conditioning on deductions greater than DKr 20,000 creates selection bias, we verified that conditioning on deductions in the previous year produces similar results ($b = 0.54$). We also ran a series of placebo tests conditioning on having deductions above thresholds ranging from -20,000 to 40,000 and found no bunching at any points in the wage earnings distribution except for the statutory kink and the pension kink.

³¹There is also no bunching ($b = -0.01$) at the middle tax pension kink (the point at which individuals who are at the pension cap begin paying the middle tax). This finding further supports prediction 1 by showing that size matters regardless of scope: a large kink that affects few workers generates more bunching than a small kink that affects few workers.

³²This endogeneity problem did not arise in testing the first two parts of prediction 2 because they did not require comparisons between individuals with different levels of deductions.

³³We exclude groups with a fraction of workers with small deductions in the bottom and top 5% of the distribution, as there are too few observations to estimate b in equal-width bins in the tails.

result confirms that tax incentives that affect a larger group of workers generate large observed elasticities. Workers with small deductions can rely on firm bunching to reach the top kink, whereas workers with large deductions need to actively search for a less common job.

Although our model does not fully characterize the dynamics of adjustment to tax changes, firm responses appear to play a central role in earnings dynamics empirically. To characterize earnings dynamics, we define an indicator for whether an individual's change in wage earnings from year t to year $t + 2$ is within DKr 7500 (the width of our bunching window) of the change in the top tax bracket cutoff from year t to year $t + 2$. This indicator measures whether an individual's earnings tracks the movement in the kink over time. Figure 17a plots the fraction of individuals who track the movement in the kink vs. the level of wage earnings in the base year relative to the statutory kink. The propensity to track the movement in the kink is highest for individuals near the kink to begin with. Figure 17b replicates Figure 17a for the pension kink, focusing on individuals with deductions greater than 20,000 in year t , as in Figure 15a. Individuals at the pension kink in year t do not have any excess propensity to track the movement in the pension kink. Instead, firm bunchers at the statutory kink (located at approximately DKr -33,000 in Figure 17b), exhibit a higher propensity to move with the kink even though they have no incentive to do so. In sum, individuals who reach the kink via firm bunching move with the kink whereas those who get there through individual job search do not. Intuitively, firms adjust the packages they offer as the aggregate distribution of workers' tax preferences change, whereas workers do not pay search costs to switch jobs and actively track the kink themselves.³⁴

We conclude that firm responses play a central role in shaping the effects of tax changes on equilibrium labor supply. Firm responses may be particularly easy to detect in Denmark because collective bargaining facilitates such responses. While collective bargaining is less common in economies such as the U.S., technological constraints force coordination of work schedules and lead to hours constraints and other institutional rigidities in all labor markets. The key lesson of the evidence here is that these constraints are endogenous to the tax regime. Although changes in taxes may not induce sharp, immediate responses by firms as in Denmark, they could induce changes in norms and job characteristics over time.

³⁴These results also provide further evidence that the difference in bunching at the top and middle kinks is not driven by heterogeneous elasticities. If individuals near the top tax cutoff were simply more elastic and did not face adjustment costs, they would track the movement of the top kink over time.

IV.C Prediction 3: Correlation Between Individual and Firm Bunching

We test the third prediction of the model by examining the correlation between individual and firm bunching across occupations. As above, we measure firm bunching b_F by the excess mass in the wage earnings distribution at the statutory top tax cutoff for individuals who have more than DKr 20,000 in deductions (and therefore have no incentive to locate at the statutory kink). We define individual bunching b_I as the excess mass at the pension kink in the wage earnings distribution for individuals with more than DKr 20,000 in deductions. Note that b_F and b_I are estimates of bunching at two *different* kinks for the same group of individuals, and thus are not mechanically related.

Figure 18 plots the estimates of b_F vs. estimates of b_I across occupations defined at the 2 digit ISCO level.³⁵ The (unweighted) correlation between b_F and b_I is 0.65 and is significantly different from 0 with $p < 0.001$. In a regression weighted by occupation size, 64% of the variation in b_F is explained by the variation in b_I . Note that the few negative point estimates of b_I and b_F are not significantly different from zero. We cannot interpret this correlation as evidence that differences in individuals' preferences *cause* changes in firm behavior as they could also be driven by sorting of workers into occupations that suit their tastes. Nevertheless, the evidence is consistent with the model's prediction that firms cater to their workers' tax-distorted preferences in equilibrium.³⁶

Differences across occupations drive much of the heterogeneity in bunching across demographic groups documented above in Figure 4. To illustrate this, Figures 19a and 19b show the distributions of taxable income for women ($b = 1.37$) and men ($b = 0.46$). The income distribution for men shown in circles in Figure 19b reweights the sample to match the observed distribution of occupations for women (following Dinardo, Fortin and Lemieux (1996)). This distribution effectively places more weight on men who work in female-dominated occupations such as teaching. This simple reweighting increases the excess mass observed among men to $b = 0.85$, closing nearly half the gap in observed elasticities between men and women.³⁷ These results underscore the importance of the constraints imposed by firms in determining workers' responses to tax policies.

³⁵Because we include only individuals with net deductions larger than DKr 20,000 to estimate b_F and b_I , there are too few observations to estimate the values at the 4 digit SIC level.

³⁶Further supporting this result, there is no firm bunching at the statutory middle kink, consistent with the lack of individual bunching at the middle tax pension kink. Individuals do not search for jobs at the middle kink because of search costs, and as a result firms do not provide such jobs in equilibrium.

³⁷The differences in elasticities between men and women may still arise from preference heterogeneity, as firms in female-dominated occupations may offer more jobs that pay salaries at the kink to cater to their employees' preferences.

IV.D Self-Employed Individuals

The self-employed are a useful comparison group because they face much smaller frictions in adjusting taxable income than wage earners. They are not subject to hours constraints imposed by firms and do not need to search for a different job to change their earnings. They can also easily change reported taxable incomes, either by shifting realized income across years or by under-reporting taxable incomes.³⁸ Therefore, we expect that the model’s three predictions should not apply to the self-employed.

Figure 20 replicates the key graphs shown above, restricting attention to the self-employed. Figure 20a shows that the self-employed exhibit extremely sharp bunching at the top kink, consistent with their ability to adjust their income more easily. The estimated excess mass is $b = 18.4$ at the top kink, dwarfing the excess mass for wage earners and implying an observed elasticity of 0.23. Figure 20b shows that unlike wage earners, the self-employed also bunch sharply at the middle tax kink. The observed elasticity at the middle kink is 0.10. We believe that the observed elasticity at the middle kink is smaller than that at the top kink because capital income is subject to the middle tax but not the top tax. Self-employed individuals are allowed to reclassify some of their profits as capital income, creating an added margin of response at the top tax cutoff. Consistent with this explanation, self-employed individuals with capital income less than DKr 1,000 in magnitude have an observed elasticity of 0.15 at the middle kink vs. 0.18 at the top kink.

Figure 20c tests for “firm bunching” by plotting the distribution of self-employment income around the statutory kink for individuals with deductions larger than DKr 20,000. Unlike wage earners, self employed individuals with large deductions exhibit no excess mass around the statutory kink. As a result, self employed individuals with common tax preferences (small deductions) bunch just as much as those with uncommon tax preferences (large deductions). This is shown in Figure 20d, which is constructed using mean group deductions in the same way as Figure 16.

These “placebo tests” confirm that our three predictions do not apply to the self-employed. We suspect that much of the bunching among the self-employed is not driven by changes in labor supply but rather by intertemporal shifting and evasion. LeMaire and Schjerning (2007) demonstrate using the same Danish data that the self-employed adjust their retained earnings and profit distributions over time to remain below the top tax threshold in each year. Kleven et al. (2009a) uncover substantial tax evasion among the self-employed and estimate that 40% of the bunching at the top kink is driven by tax evasion. Regardless of which margin the self employed use, we can

³⁸The Danish tax code allows the self-employed to shift some income across years legally.

conclude that frictions are a key determinant of observed elasticities: the size and scope of tax changes matters less for margins of behavior with low frictions (changing reported taxable income or self-employment earnings) than for margins with higher frictions (changing wage earnings).

V Calibration: A Bound on the Macro Elasticity

What do our estimates of observed elasticities tell us about the structural and macro elasticities? To answer this question, we bound ε using a partial identification approach. We show that the observed elasticities from bunching at the middle and top kinks together yield a lower bound on ε . The key intuition is that ε controls the utility loss of deviating from the optimal level of hours because it determines the concavity of the utility function (Chetty 2009a). When ε is small, the gains from searching for a level of hours close to h_i^* are very large. Hence, the model cannot generate very different observed elasticities at the middle and top kinks with small ε unless search costs are very large. Therefore, by placing an upper bound on the size of search costs, we obtain a lower bound on ε .

To formalize this logic and obtain an illustrative quantitative bound on ε , we impose some parametric structure on the general model specified in section II. We assume that the distribution of job offers conditional on search and the cost of search effort e are given by:

$$(18) \quad \begin{aligned} h'_i &= eh_i^* + (1 - e)\tilde{h}_i \\ \phi_i(e) &= \phi \cdot c_i^* \cdot e^\gamma \end{aligned}$$

where $\tilde{h}_i \sim N(h_i^*, \sigma^2)$. With this search technology, an individual draws exactly her optimal hours level h_i^* with probability e ; with probability $(1 - e)$ she draws from a normal distribution centered at h_i^* with standard deviation σ . Here, ϕ measures the cost of finding a job with h_i^* hours with certainty as a percentage of optimal consumption c_i^* , while γ controls the elasticity of search costs with respect to the precision of search. With these functional forms, the structural parameters of the model are $\omega = \{\varepsilon, \gamma, \phi, \sigma, F(\alpha_i), F_Y, p\}$. Given ω , we compute the equilibrium distribution of hours $G(h)$ and the amount of bunching implied by the model for a tax system $\{\tau_1, \tau_2, K, \zeta\}$ using numerical iteration. Details of the numerical simulation methodology are given in Appendix C.

To bound ε , we calibrate the distributional parameters of the model and set identify the remaining structural parameters. We begin by normalizing $p = w = \frac{K}{1500}$, so that individuals who

work 1500 hours, the median hours in Denmark, reach the kink.³⁹ We calibrate the distribution of shocks to non-wage income F_Y using a normal distribution with mean 0 and a standard deviation of DKr 2,746, which matches the standard deviation of net deductions conditioning on the level of lagged net deductions in the data. We cannot directly calibrate the search variance parameter σ because the search process is unobserved. We therefore estimate the σ that best fits the empirical income distributions for every combination of the remaining parameters consistent with the data. The smallest σ in this set of feasible values is $\sigma = 38,000$. Since our goal is to obtain a lower bound on ε , we fix $\sigma = 38,000$ because our estimate of ε turns out to rise with σ . We calibrate the taste distribution $F(\alpha_i)$ using two separate normal distributions whose parameters are chosen to match the empirical distributions of income away from the top and middle kinks (excluding a DKr 15,000 window around the kinks as in Figure 3) for each potential value of ε . Having calibrated these distributions, we are left with three parameters $\{\varepsilon, \gamma, \phi\}$ that we set identify using two moments: observed bunching at the middle and top kinks.

We first identify ε from the observed bunching amounts given a value of ϕ . Figure 21a shows the relationship between the structural elasticity ε and the amount of bunching at the top kink b , holding fixed all other parameters. The dashed line in Figure 21a plots the amount of bunching implied by a frictionless model without uncertainty. In the frictionless model, ε has only one effect on the amount of bunching: as ε increases, the fraction of workers who optimally relocate to the kink rises, increasing b linearly. However, with frictions, ε has a second (countervailing) effect: a larger structural elasticity implies that workers have a less concave utility function and gain less from searching for a job that places them at the kink. To see this point, note that a second-order Taylor approximation to (2) yields the following expression for the utility loss from working h_i^0 hours instead of the optimal hours h_i^* :

$$(19) \quad u_i(h_i^*) - u_i(h_i^0) \simeq \frac{1}{2} \frac{1}{\varepsilon} c_i^* \cdot (\Delta \log h_i)^2$$

where $\Delta \log h = \log h_i^0 - \log h_i^*$. It follows that the utility gain from choosing hours optimally falls with ε . Therefore, fewer workers choose to search for a new job when ε is larger, *decreasing* the observed amount of bunching at kinks.

The solid lines in Figure 21a plot the relationship between ε and the amount of bunching b that would be observed at the top kink for two values of the search cost, $\phi = 0.07$ and $\phi = 0.09$. At

³⁹We can choose p without loss of generality because a model with output price $p = 1$ and disutilities $\alpha_i = \alpha'_i/p'$ is isomorphic to one with price p' and disutilities α'_i .

low elasticities, the amount of bunching increases roughly as in the frictionless case. But at higher elasticities, the frictions begin to dominate the conventional force, and increases in ε reduce the observed amount of bunching. Comparing the two solid lines, we see that as ϕ rises, the effect of frictions grows, pulling the curve down more quickly.

The horizontal dashed line at $b = 0.81$ in Figure 21a marks the empirical estimate of bunching at the top kink for all wage earners. Because the curves in Figure 21a have an inverted-U shape when $\phi > 0$, there are two values of ε consistent with the observed amount of bunching at a kink conditional on γ and ϕ . Intuitively, a given observed elasticity can be generated by either a small structural elasticity and large gains from choosing hours optimally or a large structural elasticity and small gains from choosing hours optimally.

Next, we show that there is only one combination of ε and γ (holding fixed ϕ) that fits the observed bunching at *both* the middle and the top kinks. Figure 21b plots the b predicted by the model at the middle and top kinks in the Danish system given $\phi = 0.07$ and $\gamma = 1.05$. The empirical estimates of bunching at these two kinks (for the full population of wage earners) are marked with horizontal lines. To fit the data, the curves must each intersect the lines representing the respective observed elasticity at the same value of ε . The two lower crossing values differ significantly: $\varepsilon = 0.003$ fits the middle kink and $\varepsilon = 0.01$ fits the top kink. In contrast, the upper crossing values are identical: $\varepsilon = 0.24$ fits the estimates of b at both kinks with $\gamma = 1.05$. Only a relatively large value of ε can explain the difference in observed elasticities between the middle and top kinks because a small ε implies very large gains from search and produces $\hat{\varepsilon} = \varepsilon$ at both kinks.

There are no values other than $\varepsilon = 0.24$ and $\gamma = 1.05$ that fit the data given $\phi = 0.07$ in Figure 21b; that is, ε and γ are point identified given ϕ . To bound ε , we must therefore restrict the range of ϕ . We do so by calculating agents' average utility loss due to frictions as a fraction of consumption:

$$\delta = \int \frac{[u_i(h_i^*) - u_i(h_i)]}{c_i^*} dF(\alpha_i)$$

The utility loss δ increases with ϕ because larger search costs induce more agents to choose hours further away from their optima. For each value of δ , there is a single triplet $\{\varepsilon, \gamma, \phi\}$ that matches the observed amount of bunching at the two kinks and generates an average utility loss of δ . Figure 21c plots the estimated structural elasticity $\varepsilon(\delta)$ vs. the average utility loss δ . This curve is downward sloping because the structural elasticity ε that fits a given observed elasticity $\hat{\varepsilon}$ falls as ϕ rises, as shown in Figure 21a.

We consider $\delta = 5\%$ an upper bound on the degree of frictions: individuals are unlikely to face search costs large enough that they would tolerate a 5% consumption loss every year. With $\delta = 5\%$, a structural elasticity of $\varepsilon = 0.34$ fits the observed bunching at the middle and top kinks in Denmark. For $\varepsilon < 0.34$, there is *no* combination of parameters that fits the data and generates $\delta < 5\%$. Although the exact bound of $\varepsilon \geq 0.34$ relies on parametric assumptions, we conclude that ε is an order of magnitude larger than observed elasticities from bunching or tax reforms in our data.

Figure 21d illustrates why observed elasticities understate the structural elasticity so dramatically in our calibrated model. It plots the equilibrium earnings distribution around the top kink generated by the model with $\varepsilon = 0.34$, $\gamma = 1.015$, and $\phi = 0.049$, the parameters that fit the data at the lower bound. The figure also shows the distribution generated by a frictionless model with $\varepsilon = 0.34$ and $\phi = 0$. The simulated excess mass at the kink is an order of magnitude smaller with frictions because many of those who would optimally locate at the kink are dispersed quite diffusely around the kink. Intuitively, many individuals choose not to pay search costs to find a job at the kink because the utility gains from doing so are relatively small. For example, equation (19) implies that when $\varepsilon = 0.34$, the utility loss from ignoring the 30% reduction in the net of tax wage at the kink and working an extra $\Delta \log h_i = 0.3 \cdot 0.34$ percent is approximately 1.5% of consumption.

We calibrated the model to match prediction 1 regarding the effects of the size of kinks on observed elasticities. In the calibrated model, we also find that kinks that affect a larger fraction of individuals generate greater bunching and that a larger structural elasticity leads to higher levels of both individual and firm bunching. Hence, all three predictions that we derived analytically for special cases of the model in section II hold in the general model for empirically relevant parameters.

Micro vs. Macro Elasticities. The calibrated model can be used to predict how labor supply differs in steady state across economies with different linear tax rates τ and τ' . We find that for $\tau, \tau' \in (20\%, 80\%)$, the observed macro elasticity $\hat{\varepsilon}_M = \frac{\mathbb{E} \log h_i(\tau'_1) - \mathbb{E} \log h_i(\tau_1)}{\log(1-\tau') - \log(1-\tau)}$ differs from ε by less than 0.01. Hence, as in the special cases, the structural elasticity itself determines the observed macro elasticity irrespective of adjustment costs and firm responses. It follows that $\hat{\varepsilon}_M \geq 0.34$ given frictions of $\delta < 5\%$.

Davis and Henrekson (2005) find that an intensive-margin elasticity of $\hat{\varepsilon}_M = 0.44$ fits the observed differences in hours conditional on working across OECD countries with different tax rates. Our calibration results show that such a macro elasticity could be consistent with near-zero

micro estimates of observed elasticities. Although macro elasticity estimates suffer from many omitted variable and endogeneity problems, we conclude that at least part of the discrepancy between micro and macro elasticities could be due to adjustment costs and hours constraints.⁴⁰

VI Conclusion

This paper has shown that the effects of tax policies on labor supply are shaped by adjustment costs and hours constraints endogenously chosen by firms. Because of these forces, modern microeconomic methods of estimating elasticities – focusing on policy changes that affect a subgroup of workers – may substantially underestimate the “structural” elasticities that control steady-state behavioral responses. Calibrating a model of labor supply with search costs and hours constraints to match the empirical evidence, we obtain a lower bound on the macro elasticity of $\varepsilon \geq 0.34$, an order of magnitude larger than micro estimates for wage earners in our data.

In future work, it would be useful to extend the static analysis here to a dynamic equilibrium model with frictions. The evidence on salaries tracking the movement in the kink in Figure 17 is one of many potential facts that could be used to develop and test such a model. Workers should make discrete, lumpy adjustments in hours in response to changes in tax rates, as in an (S,s) model.⁴¹ One could also explore how the dynamics of adjustment vary across firms. Firm responses should matter more in industries with greater complementarity between workers in their production technologies. In a dynamic environment with frictions, taxes may also distort other margins not considered in this paper. Firms may offer a flatter wage profile over the lifecycle to avoid pushing workers into the highest tax brackets. And workers might change human capital choices in the long run, further attenuating short run elasticity estimates.

It would also be interesting to explore the normative implications of adjustment costs and firm responses using the model proposed here. For example, the efficiency cost of a tax levied on one group of workers may depend not just upon their elasticities but also upon those of their co-workers if firms are constrained to offer similar packages to different workers. Another example concerns the prediction that it is optimal to levy higher tax rates on men than women because they are less elastic (Boskin and Sheshinski 1983, Alesina et al. 2007, Kleven et al. 2009b). If the difference in

⁴⁰Our analysis also implies an intensive-margin Frisch elasticity above 0.34, as the difference between Frisch and Hicksian elasticities is negligible with time separable utility and a non-negative uncompensated wage elasticity (Chetty 2009a).

⁴¹Testing this prediction will require hours data. We do not observe clear evidence of lumpy adjustment in the earnings measures ($w_i h_i$) studied above, presumably because the wage rate w_i changes across periods while hours h_i change infrequently.

observed elasticities across genders is caused by heterogeneity in occupational frictions rather than tastes, there may be less justification for higher taxes on secondary earners in steady state.

Finally, the results here call into question the modern empirical paradigm of using quasi-experiments that apply to specific subgroups to learn about the effects of economic policies and shocks on behavior. In settings with rigid institutional structures and frictions in adjustment, the steady-state effects of policies implemented at an economy-wide level could differ substantially from the effects of such experiments. Using large, broad policy variation to partially identify a structural model may be a promising approach to learn about steady-state effects in many applications beyond taxation and labor supply.

References

1. Altonji, Joseph G. and Christina H. Paxson (1988). "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs," *Journal of Labor Economics*, 6(2): 254-76.
2. Alesina, Alberto F., Edward L. Glaeser, and Bruce Sacerdote (2005). "Work and Leisure in the U.S. and Europe: Why So Different?" *NBER Maco Annual*, pp. 1-99.
3. Alesina, Alberto F., Andrea Ichino, and Loukas Karabarbounis (2007). "Gender-Based Taxation and the Division of Family Chores." NBER Working Paper 13638.
4. Andersen, Torben M. and Michael Svarer (2007). "Flexicurity: Labour Market Performance in Denmark," *CESifo Economic Studies*, Oxford University Press, 53(3): 389-429
5. Blundell, Richard, and Thomas MaCurdy (1999). "Labor Supply: A Review of Alternative Approaches." *Handbook of Labor Economics*, Vol 3A, Ashenfelter and Card eds., pp. 1559-1695.
6. Boskin, Michael J. and Eytan Sheshinski (1983). "Optimal Tax Treatment of the Family: Married Couples," *Journal of Public Economics* 20, pp. 281-297.
7. Chetty, Raj (2009a). "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." NBER Working Paper No. 15616.
8. Chetty, Raj (2009b). "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy* 1 (2): 31-52.
9. Chetty, Raj and Emmanuel Saez (2009). "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients." NBER Working Paper No. 14836.
10. Chetty, Raj, Anders Frederiksen, John Friedman, Tore Olsen, and Luigi Pistaferri (2010). "What Determines Knowledge of Tax Policies? Evidence from Denmark." Harvard University Working Paper.
11. Cogan, John (1981), "Fixed Costs and Labor Supply." *Econometrica*, Vol. 49(4): 945-63.

12. Davis, Steven and Magnus Henrekson (2005). "Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich-Country Comparisons," in *Labour Supply and Incentives to Work in Europe*, eds, R. Gomez Salvador, A. Lamo, B. Petrongolo, M. Ward, and E. Wasmer, Edward Elgar Press.
13. Dickens, William T. and Shelly J. Lundberg (1993). "Hours Restrictions and Labor Supply." *International Economic Review*, 34(1): 169-92.
14. DiNardo, John, Nicole M. Fortin, Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach." *Econometrica*, Vol 64, pp. 1001-1044.
15. Feldstein, Martin (1999). "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics*, Vol. 81(4), pp. 674-680.
16. Goolsbee, Austin (2000). "What Happens When You Tax the Rich? Evidence from Executive Compensation." *Journal of Political Economy*, Vol 108, pp. 352-378.
17. Gruber, Jonathan and Emmanuel Saez (2002). "The Elasticity of Taxable Income: Evidence and Implications" *Journal of Public Economics*, Vol. 84(1), pp. 1-32.
18. Ham, J.C. (1982). "Estimation of a Labour Supply Model with Censoring Due to Unemployment and Underemployment." *Review of Economic Studies*, Vol. 49(3), pp. 335-354.
19. Hausman, Jerry A. (1981). "Labor Supply." In *How Taxes Affect Economic Behavior*, eds. Henry Aaron and Joseph Pechman.
20. Kleven, Henrik and Claus Kreiner (2006). "Welfare and Taxation" in *Tax, Work, and Equality: An Investigation of the Danish Tax and Welfare System*. Copenhagen: Gyldendal.
21. Kleven, Henrik, Martin Knudsen, Claus T. Kreiner, Søren Pedersen, and Emmanuel Saez (2010). "Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark." NBER Working Paper 15769.
22. Kleven, Henrik, Claus Kreiner, and Emmanuel Saez (2009b). "The Optimal Income Taxation of Couples." *Econometrica*, Vol 77, pp. 537-560.
23. Kleven, Henrik and Esben A. Schultz (2010). "Estimating Taxable Income Responses Using Danish Tax Reforms." LSE Working Paper.

24. LeMaire, Daniel, and Bertel Schjerning (2007). "Earnings, Uncertainty, and the Self-Employment Choice." CBER Discussion Paper DP 2007-04.
25. Lindbeck, Assar (1995). "Welfare State Disincentives with Endogenous Habits and Norms." *The Scandinavian Journal of Economics*, Vol. 97(4), pp. 477-494.
26. MaCurdy, Thomas, David Green and Harry Paarsche (1990). "Assessing Empirical Approaches for Analyzing Taxes and Labor Supply." *The Journal of Human Resources*, Vol 25, pp. 415-490.
27. Ministry of Taxation (2002). *Skatteberegningssreglerne, før og nu*.
http://www.skm.dk/public/dokumenter/Tal45/Skatteberegning90/Skatteberegningssreglerne_2.pdf
28. Moffitt, Robert A. (1990). "The Econometrics of Kinked Budget Constraints." *Journal of Economic Perspectives*, Vol 4, pp. 119-139.
29. OECD (2009). Factbook 2009: Economic, Environmental and Social Statistics.
<http://puck.sourceoecd.org/vl=13806301/cl=14/nw=1/rpsv/factbook2009/06/03/02/index.htm>
30. Rogerson, Richard and Johanna Wallenius (2009). "Micro and Macro Elasticities in a Life Cycle Model with Taxes." Forthcoming, *Journal of Economic Theory*.
31. Rosen, Harvey S. (1976). "Taxes in a Labor Supply Model with Joint Wage-Hours Determination," *Econometrica*, 44(3): 485-507.
32. Saez, Emmanuel (2009). "Do Taxpayers Bunch at Kink Points?" Forthcoming. *American Economic Journal: Economic Policy*.
33. Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz (2009). "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." Forthcoming, *Journal of Economic Literature*.
34. Slemrod, Joel B. (1995). "Income Creation or Income Shifting? Behavioral Responses to the Tax Reform Act of 1986." *American Economic Review*, Vol. 85(2): pp. 175-80
35. Slemrod, Joel B. (1998). "Methodological Issues in Measuring and Interpreting Taxable Income Elasticities." *National Tax Journal*, Vol. 51(4), pp. 773-788.
36. Slemrod, Joel B. and Shlomo Yitzhaki (2002). "Tax avoidance, evasion, and administration." In *Handbook of Public Economics*, Alan Auerbach and Martin Feldstein (eds.), pp. 1423-1470.

Appendix A: Theoretical Derivations

Predictions 1-3 for Tax Reforms. We introduce a second period in the model to analyze the effects of tax reforms. At the beginning of the second period, the government announces an unexpected tax reform that raises the linear tax rate for workers of type $s_i = L$ from τ to $\tilde{\tau}$. Let $\Delta\tau = \tilde{\tau} - \tau$ and $\Delta\log(1 - \tau) = \log(1 - \tilde{\tau}) - \log(1 - \tau)$.

We model the search process in period 1 exactly as above. Because the tax reform is unanticipated, worker and firm behavior in period 1 is the same as in the static model. In period 2, firms can change the hours they require from workers at no cost and again choose h_j to maximize profits $\pi = ph_j - w(h)h_j$. Free entry implies that the equilibrium wage rate is $w(h) = w$ in period 2 as in period 1. Workers associated with a firm that changes its hours requirement are forced to work that new level of hours unless they switch jobs. After seeing the full distribution of firm responses in period 2, workers can pay a search cost $\phi_i(e)$ to find a new job, where the precision of the search related inversely to $e \in [0, 1]$ as in the static model. Equilibrium in the second period requires that the aggregate distribution of jobs posted by firms in period 2 coincides with those chosen by workers who maximize utility net of search costs. A full characterization of dynamics requires assumptions about the specific firms that move in order to shift the old equilibrium distribution of jobs to the new equilibrium distribution. The results we derive below rely only on aggregate dynamics and thus do not require such assumptions.

Let h_{it}^* denote worker i 's optimal labor supply choice in period t and h_{it} her actual choice in equilibrium. We characterize the observed elasticity from the tax reform $\hat{\varepsilon}_{TR} = \frac{\mathbb{E} \log h_{i2} - \mathbb{E} \log h_{i1}}{\Delta \log(1 - \tau)}$ in each of the special cases analyzed in Section II in turn.

Special Case 1. In the frictionless benchmark model, $\phi = 0$ for all workers, in which case workers set $h_{it}^*(\tau) = \alpha_i (w(1 - \tau))^\varepsilon$ in both periods. It follows immediately that the observed elasticity from a tax reform $\hat{\varepsilon}_{TR} = \frac{\mathbb{E} \log h_{i2} - \mathbb{E} \log h_{i1}}{\Delta \log(1 - \tau)} = \varepsilon$.

Special Case 2. In the second special case, $\phi_i(e) = \phi$ is constant and a measure zero set of workers faces the linear tax schedule ($\zeta = 1$), so the equilibrium distribution of hours $G(h)$ is unchanged across the two periods. In the second period, a worker's first-period job h_{i1} functions as an initial offer, just as the initial draw h_i^0 did in the first period. A worker pays to switch to his optimal job h_{i2}^* iff $h_{i1} \notin [\underline{h}_{i2}, \bar{h}_{i2}]$, where the thresholds are defined as in equations (7) and (8). When $\Delta\tau = 0$, the new bounds coincide with the old: $\bar{h}_{i2} = \bar{h}_{i1}$ and $\underline{h}_{i2} = \underline{h}_{i1}$. As the size of the tax reform grows, more workers have $h_{i1} \notin [\underline{h}_{i2}, \bar{h}_{i2}]$ because $\frac{\partial \underline{h}_{i2}}{\partial \Delta\tau} < 0$ and $\frac{\partial \bar{h}_{i2}}{\partial \Delta\tau} < 0$. Therefore the

fraction of workers paying to search increases. Average labor supply for those with $s_i = L$ in the second period can be written as

$$(20) \quad h_2 = \int \left[q_{i2} h_{i2}^* + \int_{\underline{h}_{i1}}^{\bar{h}_{i2}} h dG(h) \right] dF(a_i)$$

where $q_{i2} = 1 - G(\bar{h}_{i2}) + G(\underline{h}_{i1})$ is the fraction of workers that switch jobs after the reform.

As the size of the tax reform grows large, the observed elasticity converges to ε : $\lim_{\Delta\tau \rightarrow \infty} \hat{\varepsilon}_{TR} = \varepsilon$. Intuitively, for a sufficiently large tax reform, $\bar{h}_{i2} \leq \underline{h}_{i1}$, in which case all workers pay to search ($q_{i2} = 1$) and set $h_{i2} = h_{i2}^*$. Although workers do not all have $h_{i1} = h_{i1}^*$, the change in average hours grows large relative to $h_{i1} - h_{i1}^*$ as $\Delta\tau \rightarrow \infty$, and thus $\hat{\varepsilon}_{TR} \rightarrow \varepsilon$. While $\hat{\varepsilon}_{TR}$ always converges to ε , the derivative $\frac{\partial \hat{\varepsilon}_{TR}}{\partial \Delta\tau}$ can only be signed by making assumptions about the job offer distribution $G(h)$. Suppose that the distribution of preferences are such that the equilibrium distribution of jobs $G(h)$ is uniform for those with $s_i = NL$, who do *not* face the tax reform. Under this assumption, the fraction of workers who reoptimize following the tax change q_{i2} increases monotonically from 0 to 1 as the size of the reform increases and hence $\frac{\partial \hat{\varepsilon}_{TR}}{\partial \Delta\tau} > 0$.

Combining these results yields a prediction for tax reforms analogous to Prediction 1.

PREDICTION A1: *When workers face search costs,*

(a) *the observed elasticity from tax reforms converges to ε as the size of the tax change grows:*

$$\lim_{\Delta \log(1-\tau) \rightarrow \infty} \hat{\varepsilon}_{TR} = \varepsilon$$

(b) *If the offer distribution $G(h)$ is uniform, $\hat{\varepsilon}$ rises with $\Delta\tau$:*

$$\partial \hat{\varepsilon}_{TR} / \partial \ln(1-\tau) > 0$$

Special Case 3. In the third special case, $\phi_i(e) = 0$ for a fraction δ of workers and $\phi_i(e) = \infty$ for the rest, and $\zeta \in (0, 1)$. In both periods, the equilibrium distribution of hours $G(h(\tau)) = G^*(h(\tau))$, the optimal distribution of hours, following the same logic as in the text. Let $\Delta \log h_L^* = \varepsilon \cdot (\log(1 - \tilde{\tau}) - \log(1 - \tau))$ denote the optimal change in hours for those facing the linear tax. The actual change in hours for this group is $\Delta \log h_L = (\delta + (1 - \delta)(1 - \zeta)) \Delta \log h_L^*$. The first term in this expression is the individual response (the analog of individual bunching), $\Delta \log h_I = \delta \Delta \log h_L^*$. The second term is the firm response (the analog of firm bunching), $\Delta \log h_F = (1 - \delta)(1 - \zeta) \Delta \log h_L^*$.

The change in hours for those with $s_i = NL$ is $\Delta \log h_{NL} = (1 - \delta)(1 - \zeta) \Delta \log h_L^* = \Delta \log h_F$, providing an empirical measure of the firm response. Recognizing that the observed elasticity is

$$\hat{\varepsilon}_{TR} = \frac{\Delta \log h_L}{\Delta \log(1 - \tau)} = (\delta + (1 - \delta)(1 - \zeta)) \varepsilon$$

the analogs of predictions 2 and 3 follow immediately.

PREDICTION A2: *Search costs interact with hours constraints to generate firm responses to tax reforms. The size of the firm response and observed elasticity rises with the fraction of workers who face a tax reform:*

$$\begin{aligned} \Delta \log h_F &= \Delta \log h_{NL} > 0 \text{ iff } \zeta < 1 \\ \frac{\partial \Delta \log h_F}{\partial(1 - \zeta)} &> 0, \quad \frac{\partial \hat{\varepsilon}_{TR}}{\partial(1 - \zeta)} > 0. \end{aligned}$$

PREDICTION A3: *Firm and individual responses to a tax reform are positively correlated across occupations:*

$$\text{cov}(\Delta \log h_I^q, \Delta \log h_F^q) > 0$$

Non-Constant Structural Elasticities. Suppose agents have quasi-linear utilities of the form $u_i(c, h) = c - \frac{1}{\alpha_i} \psi(h)$. This utility permits the structural elasticity of labor supply $\varepsilon = \frac{\partial \ln h^*}{\partial \ln((1 - \tau_1)w)}$ to vary arbitrarily with the net-of-tax rate depending upon $\psi''(h)$. In the frictionless model, workers who face an increase in their marginal tax rates from τ_1 to τ_2 at an earnings level of K bunch at the kink iff $\alpha_i \in [\underline{\alpha}(\tau_1), \bar{\alpha}(\tau_2)]$, where $\underline{\alpha}(\tau_1) = \psi'(h_K)/((1 - \tau_1)w)$ and $\bar{\alpha}(\tau_2) = \psi'(h_K)/((1 - \tau_2)w)$. The amount of bunching at the kink is therefore $B^*(\tau_1, \tau_2) = \int_{\underline{\alpha}(\tau_1)}^{\bar{\alpha}(\tau_2)} dF(\alpha_i)$. It follows that for any tax rates $\tau_1 < \tau_2 < \tau_3$, the amount of bunching created from two smaller kinks is exactly equal to the bunching created at one larger kink:

$$B^*(\tau_1, \tau_3) = B^*(\tau_1, \tau_2) + B^*(\tau_2, \tau_3)$$

Now consider special case 2 of the model with frictions, where agents pay a fixed cost ϕ to search. Here, the amount of bunching is

$$B(\tau_1, \tau_2) = \theta(\tau_2 - \tau_1) \int_{\underline{\alpha}(\tau_1)}^{\bar{\alpha}(\tau_2)} dF(\alpha_i)$$

where the fraction of workers who pay the search cost to locate at the kink (θ) increases with the change in tax rates at the kink ($\tau_2 - \tau_1$). Therefore the model with frictions instead implies that bunching at one large kink is greater than the sum of bunching at two smaller kinks:

$$B(\tau_1, \tau_3) > B(\tau_1, \tau_2) + B(\tau_2, \tau_3).$$

Appendix B: Data

We merge selected variables from the following registers available at the Center for Applied Microeconometrics at University of Copenhagen through Statistics Denmark: a) the Income Statistics Register, which covers everyone who is tax liable in Denmark, b) the Population Register, which covers the entire population on December 31st of a given year and provides basic demographic information such as age and gender, and c) the Integrated Database for Labour Market Research (IDA), which contains information on labor market experience, occupation, employment status, education, family status, etc. For every gender-age cell of the individuals between the ages of 16 and 70, we have tax records for between 99.96 and 100% of the population. We do not have tax records for people over 70 years of age, and 83% of 15 year olds have records in the tax register.

Statistics Denmark’s Employment Classification Module combines several administrative records to assign every observation in the IDA database one of eight employment codes, contained in the variable *beskst* (employment status). The employment status code distinguishes individuals who are wage earners, wage earners with unemployment income, wage earners with self employment income, and five categories of non-wage earners (self-employed, pensioners, etc.). To form our primary analysis dataset, we keep only individuals with *beskst*=4, thereby excluding all non wage earners, wage earners with self employment income, and wage earners with unemployment. Broadening this definition to include all non-self employment categories (*beskst*=4,5,7, or 8) does not affect the results; for instance, we find excess mass at the top kink of $b = 0.83$ in the broader sample compared with $b = 0.81$ for the narrower sample used in Figure 3.

To calculate marginal tax rates and income relative to the tax bracket cutoffs, we develop a tax simulator for Denmark analogous to NBER TAXSIM. Denmark has essentially an individual tax system, but there are some joint aspects, so the tax simulator uses as inputs both income related to the social security number associated with a given tax record (*pnr*) as well as that of the spouse for tax purposes (*henv*). The municipality of residence in the previous year (*glskkmnr*) is used

to determine what tax rates the individual faces. For the tax payer and his or her spouse, the variables used in the tax simulator are primarily the personal exemption ($pfrdst$, $berfrdst$), personal income ($perindkp$, $berpi$), capital income ($kapindkp$, $berkap$), special deductions ($lignfrdp$). We also make use of some other more disaggregated variables in the tax records to account for transitional schemes and special adjustments to the tax bases. These include deductions in personal income for individual contributions to pension schemes ($kappens$, $fosfufrd$), employer contributions to capital pension schemes ($arbpen14$, $arbpen15$), and alimony paid ($underhol$).

We assess the accuracy of the tax calculator using data from the tax register on the exact amount of municipal, regional, bottom, middle and top tax paid by each individual. Our tax calculator is correct to within +/- 5DKr (\$1) of the actual amount paid for all of these taxes for 95% of the observations in the data. It is accurate to within +/-1,000DKr (\$167) for 98% of the observations. The discrepancies arise from our inability to fully model complex capital income transfer rules that apply to some spouses as well as unusual circumstances such as individuals who die during the year or those working both in Denmark and abroad who are subject to special tax treaties. Since we do not have tax records for people aged less than 15 or more than 70, we also cannot fully account for the joint aspects of the tax system for people with spouses aged less than 15 or over 70.

In addition to the variables described above used to compute taxable income and pension contributions, we also use the following source variables in our empirical analysis: wage earnings ($qlontmp2$), self-employment profits and retained earnings ($qovskvir$, $virjordind$), labor market experience ($erhver$, $erhver79$), and occupational code ($discok$). We define an individual's net deductions in the top tax base as the level of wage earnings he/she would need to start paying the top tax minus the statutory top tax cutoff (i.e. the level of total personal income at which individuals must start paying the top tax).

The STATA code and tax simulator are available from the authors and have been posted on the servers at the Center for Applied Microeconometrics.

Appendix C: Calibration Methodology

This appendix describes how we simulate the general model and bound ε to match the empirical estimates.

Numerical Simulation. Given a vector of structural parameters $\omega = \{\varepsilon, \gamma, \phi, \sigma, p, F(\alpha_i), F_Y\}$ and a tax system $\{\tau_1, \tau_2, K, \zeta\}$, we numerically simulate the equilibrium hours distribution and the

amount of bunching in the general model using iteration. We discretize the hours distribution $G(h)$ so that the distribution of earnings is evenly spaced in 500 DKr steps. Let g_j denote the fraction of agents who work hours h_j in the equilibrium defined in section II and g_j^* denote the unconstrained optimal distribution of hours choices in the frictionless model. We similarly discretize the taste distribution $F(\alpha_i)$. Let P_i denote the fraction of agents with disutility α_i .

We first solve for the worker's optimal hours choice h_i^* by numerically integrating over the distribution of unearned income $F_Y(y_i)$ to calculate each agent's expected utility. We then start from an initial offer distribution that matches the frictionless optimum $g_j^0 = g_j^*$ and compute the distribution of jobs chosen by workers after the search process (again using numerical integration over F_Y). We label the resulting distribution g_j^1 and proceed to iterate this procedure until we obtain convergence: $\sum_j (g_j^i - g_j^{i+1})^2 < 10^{-20}$. Finally, we add the distribution of non-wage income F_Y to the equilibrium earnings distribution to calculate the equilibrium distribution of taxable income. We verify that the equilibrium distribution g_j is unique by varying the initial distribution g_j^0 . Finally, we use the method in section II.B to calculate the amount of bunching at the kink (b) and apply equation (6) to recover the observed elasticity implied by this value of b .

Calibration. We specify the tax system parameters using the values for the year 2000, shown in Figure 2a. The top tax parameters are $\{\tau_1 = 0.55, \tau_2 = 0.68, K = 267,600, \zeta = 0.6\}$ and the middle tax as $\{\tau_1 = 0.49, \tau_2 = 0.55, K = 164,300, \zeta = 0.4\}$. Given each value of ε , we first estimate the parameters $\{\phi, \gamma, \sigma, F(\alpha_i)\}$ by minimizing the sum of squared errors between the simulated and empirical distributions of taxable income around the two kinks. We then fix the standard deviation at the lowest value of σ that fits the data for any value of ε :

$$\sigma = \min \sigma^*(\varepsilon) \text{ s.t. } b^T = \widehat{b}^T, b^M = \widehat{b}^M.$$

Holding fixed σ at this lower bound, we find the values of ϕ and γ that generate simulated amounts of bunching equal to the empirical estimates of b for all wage earners, re-estimating the best fit $F(\alpha_i)$ distribution for each ε . This procedure ensures that the lower bound on ε is identified solely from the amount of bunching at the two kinks. Finally, we calculate the δ implied by the parameters associated with each value of ε by calculating the average value of $u(h_i^*) - u(h_i)$ for agents with observed earnings with DKr 50,000 of the top kink.

TABLE 1

Summary Statistics for the Danish Population and Estimation Samples, 1994-2001

	Population	Wage Earners	Wage Earners < DKr 50,000 from top tax cutoff	Self employed
	(1)	(2)	(3)	(4)
<u>Demographics:</u>				
Age	40.91	39.17	41.43	46.02
Children	0.62	0.67	0.71	0.70
Labor market experience (years)	12.46	15.42	18.77	9.46
College education	17.61%	22.76%	28.54%	17.74%
Female	49.61%	48.17%	39.17%	24.40%
Married	50.62%	53.64%	58.68%	67.34%
<u>Income:</u>				
Wage earnings	149,254	236,478	269,340	38,343
Other personal income	42,642	9,408	2,747	153,467
Total personal Income	180,213	227,359	251,145	188,854
Net capital income	-10,672	-15,819	-19,570	-7,785
<u>Deductions:</u>				
Net deductions	-40,687	-13,151	-6,381	-31,996
Net deductions <7,500	43.25%	59.36%	69.11%	23.84%
Net deductions-Pension kink <7,500	2.03%	2.72%	2.96%	5.07%
Individual pension contributions	4,316	4,217	4,535	16,709
Employer pension contributions	7,584	12,065	13,131	2,123
<u>Tax Payments:</u>				
Predicted liability accurate within 5 DKr	95.11%	94.83%	94.47%	93.62%
Pays the middle tax	50.38%	74.23%	95.57%	45.48%
Pays the top tax	18.06%	25.87%	33.53%	23.61%
2-year growth in net-of-tax rate (NTR)	1.68%	2.25%	2.25%	1.07%
Std dev of 2-year growth in NTR	4.50%	4.95%	4.95%	6.80%
Number of observations	30,492,819	17,866,090	6,788,235	1,846,064

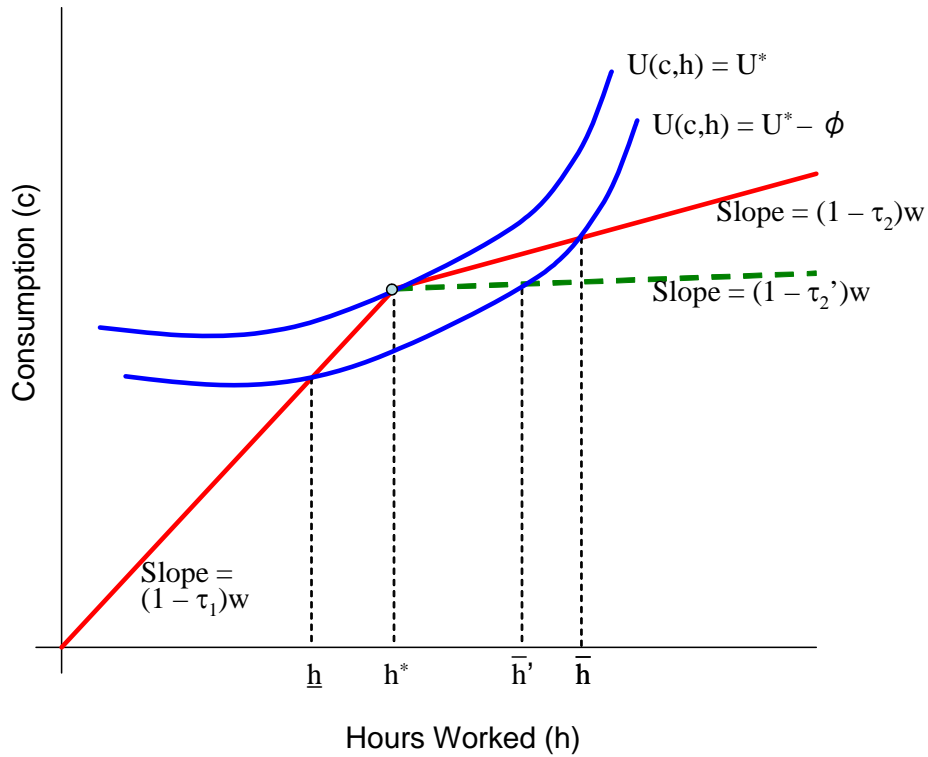
NOTE--Table entries are means unless otherwise noted. Column 1 is based on the full population of Denmark between ages 15-70 from 1994-2001. Column 2 includes all wage earners, the primary estimation sample. Column 3 includes only the subset of wage earners for whom |taxable income - top tax cutoff| < 50,000, i.e. the individuals in Figure 3. Column 4 considers individuals who report positive self-employment income. All monetary values are in real 2000 Danish Kroner. Children are the number of children younger than 18 living with the individual. Personal income refers to all non-capital income. Net capital income refers to capital income minus payments such as mortgage interest. Net deductions refer to deductions from the top tax base such as individual pension contributions minus non-wage income such as taxable gifts. Net of tax rate is one minus the marginal tax rate predicted by our tax simulator.

TABLE 2
Observed Elasticity Estimates using Small Tax Reforms

<i>Dependent Variable: Log Change in Wage Earnings</i>						
Subgroup:	All Wage Earners Baseline	Full Controls	Married Females	High-Experience Married Female Professionals	Wage Earners > 200K	Married Females > 200K
	(1)	(2)	(3)	(4)	(5)	(6)
log change in net-of-tax rate ($\Delta \log(1-t)$)	-0.005 (0.003)	-0.007 (0.004)	0.002 (0.005)	0.001 (0.011)	-0.001 (0.003)	0.013 (0.006)
Labor income spline	x	x	x	x	x	x
Personal income spline	x		x	x	x	x
Year fixed effects	x	x	x	x	x	x
Age fixed effects	x	x	x	x	x	x
Region fixed effects		x				
Occupation fixed effects		x				
Capital income spline		x				
Gender, Marital status		x				
Observations	11,512,625	8,189,920	3,136,894	156,527	7,480,900	1,767,737

Notes: Standard errors clustered by individual reported in parentheses. Dependent variable in all specifications is nominal two-year growth rate in wage earnings. Independent variable of interest is two-year growth rate in net-of-tax rate, instrumented using two-year growth rate in simulated net-of-tax rate using base-year variables. Coefficients reported can be interpreted as observed wage earnings elasticities from tax reforms. All specifications include 10-piece wage earnings and total personal income splines as well as age and year fixed effects. Column 2 also includes a 10 piece capital income spline, gender and marital status indicators, and region and occupation fixed effects. Occupation fixed effects are available only for a subset of years and observations. Column 4 restricts attention to married female professionals with more than 19 years of labor market experience. Columns 5 and 6 restrict attention to individuals with more than DKr 200,000 of wage earnings in the base year.

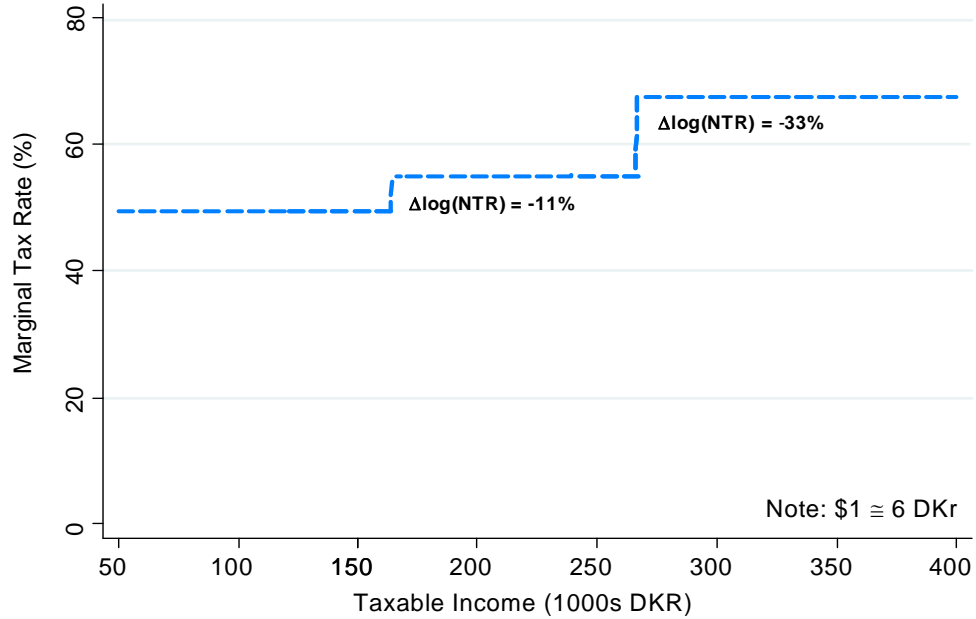
FIGURE 1
Bunching at Kinks with Search Costs



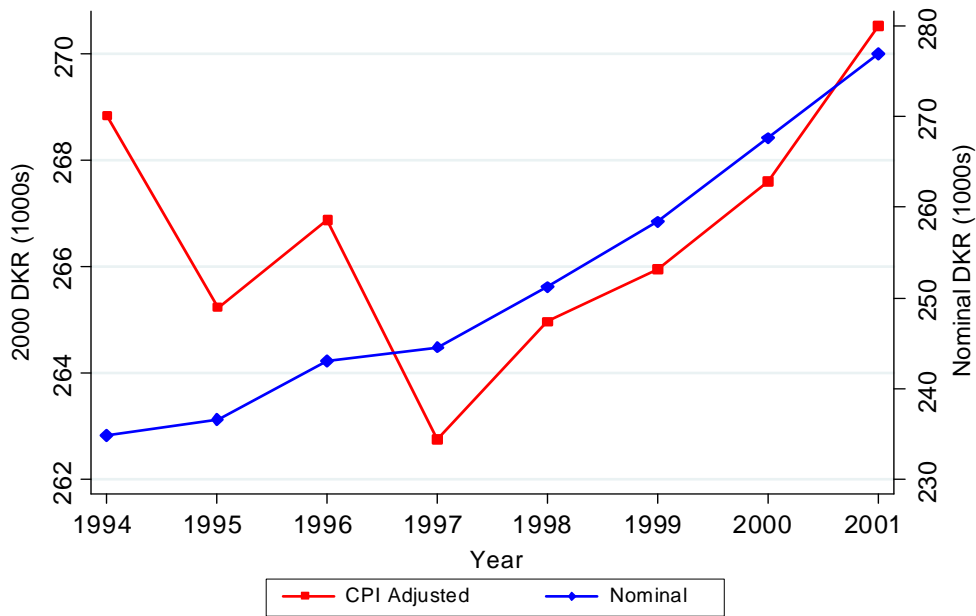
Notes: This figure illustrates how search costs affect bunching at kinks. The two-bracket tax system creates the kinked budget set shown in red. The worker's indifference curves are shown by the blue isoquants. This worker's optimal labor supply is to set $h^* = h_K$, placing her at the kink. The lower indifference curve shows the optimal utility minus the search cost ϕ . If the worker draws an initial hours offer between \bar{h} and \underline{h} , she will not pay ϕ to relocate to the kink. As the tax change at the bracket cutoff increases in magnitude (shown by the green dashed line), the inaction region shrinks to $(\underline{h}, \bar{h}')$, leading to a larger observed elasticity from bunching.

FIGURE 2
The Danish Income Tax System

(a) Marginal Tax Rates in Denmark in 2000

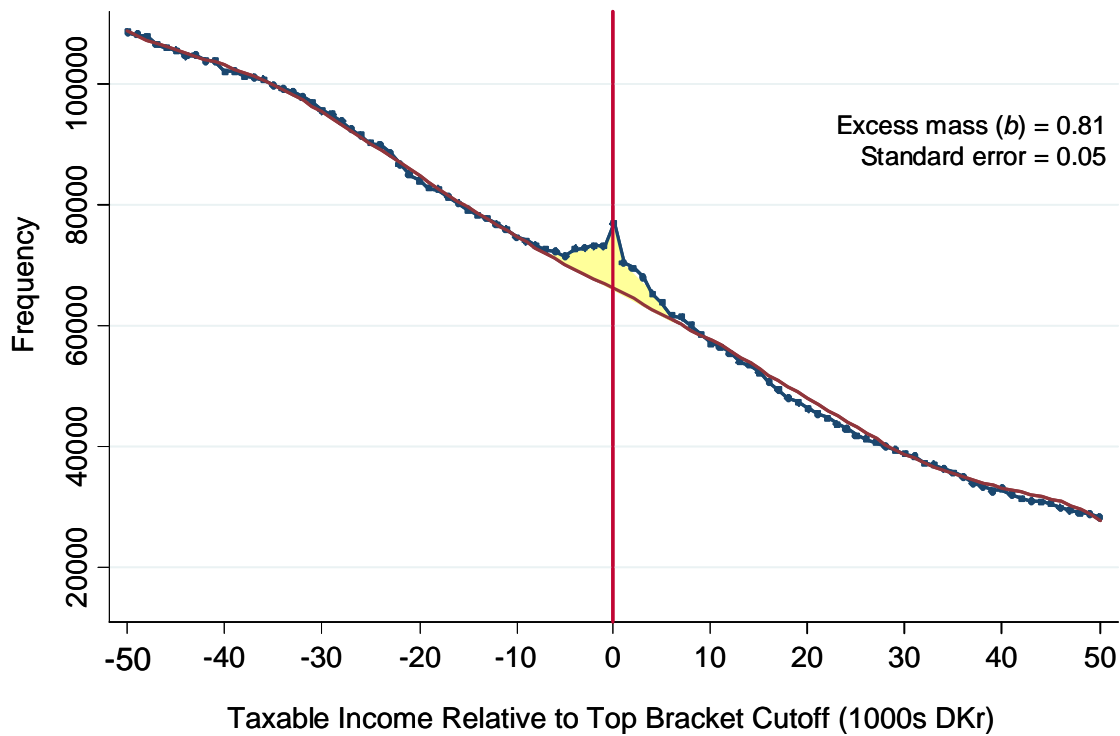


(b) Movement in Top Tax Cutoff Over Time



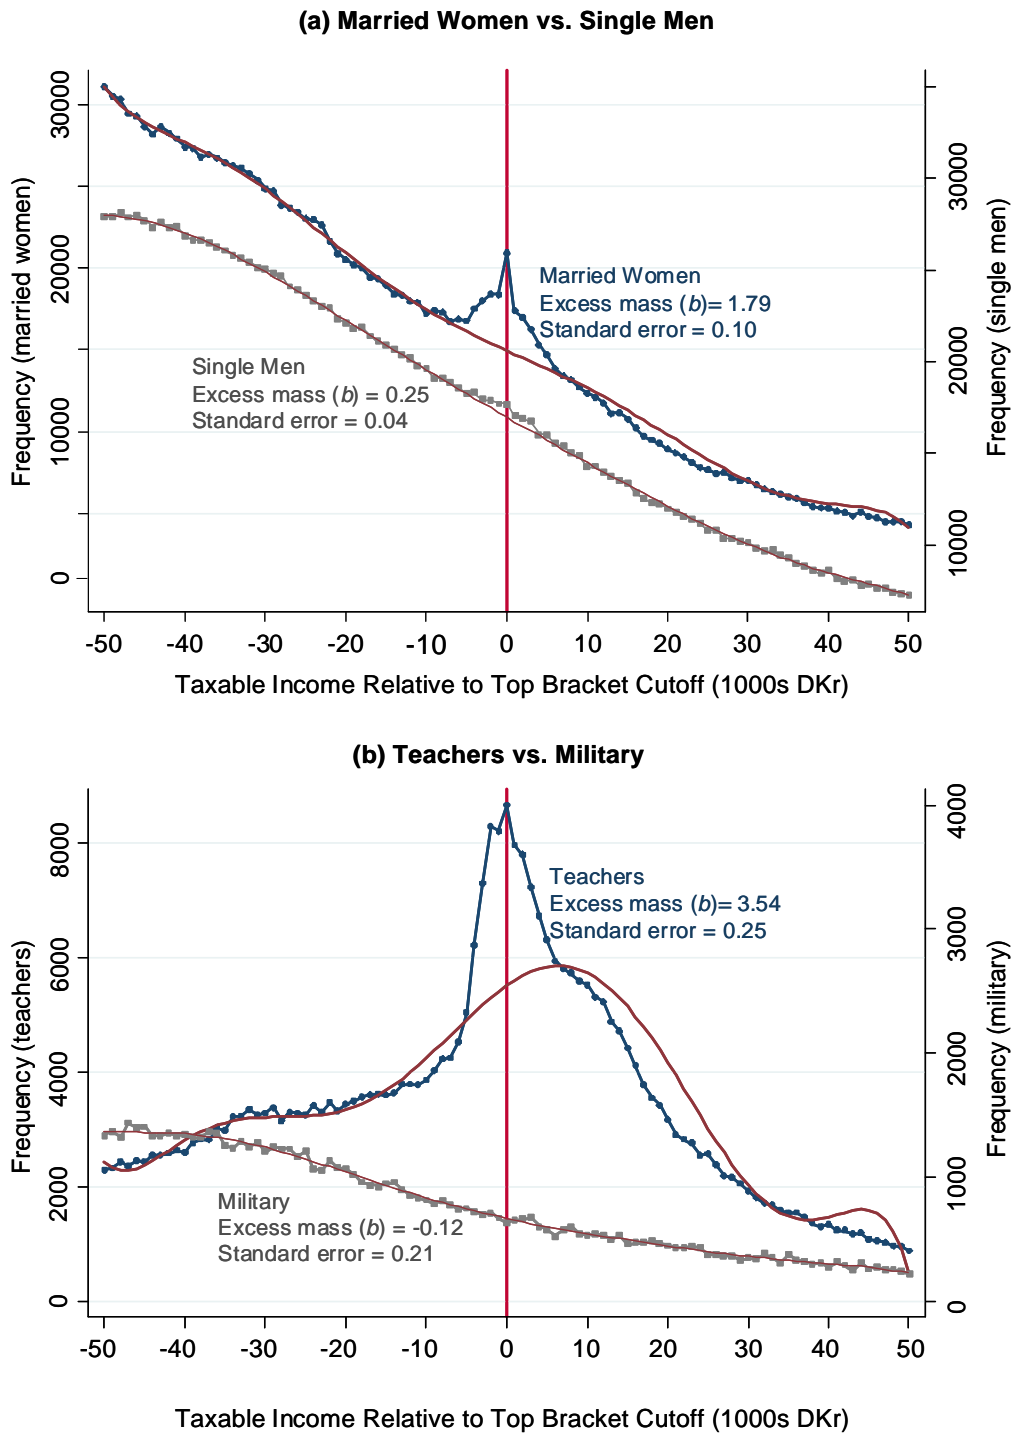
Notes: Panel (a) plots the marginal tax rate in 2000 vs. income in Denmark, including the national tax, regional tax, and average municipal tax. Panel (b) plots the level of taxable income above which earners must pay the top bracket national tax. The series in blue diamonds, plotted on the right y-axis, shows the nominal cutoff; the series in red squares, plotted on the left y-axis, shows the cutoff in real 2000 DKr.

FIGURE 3
Income Distribution around the Top Tax Cutoff for Wage Earners



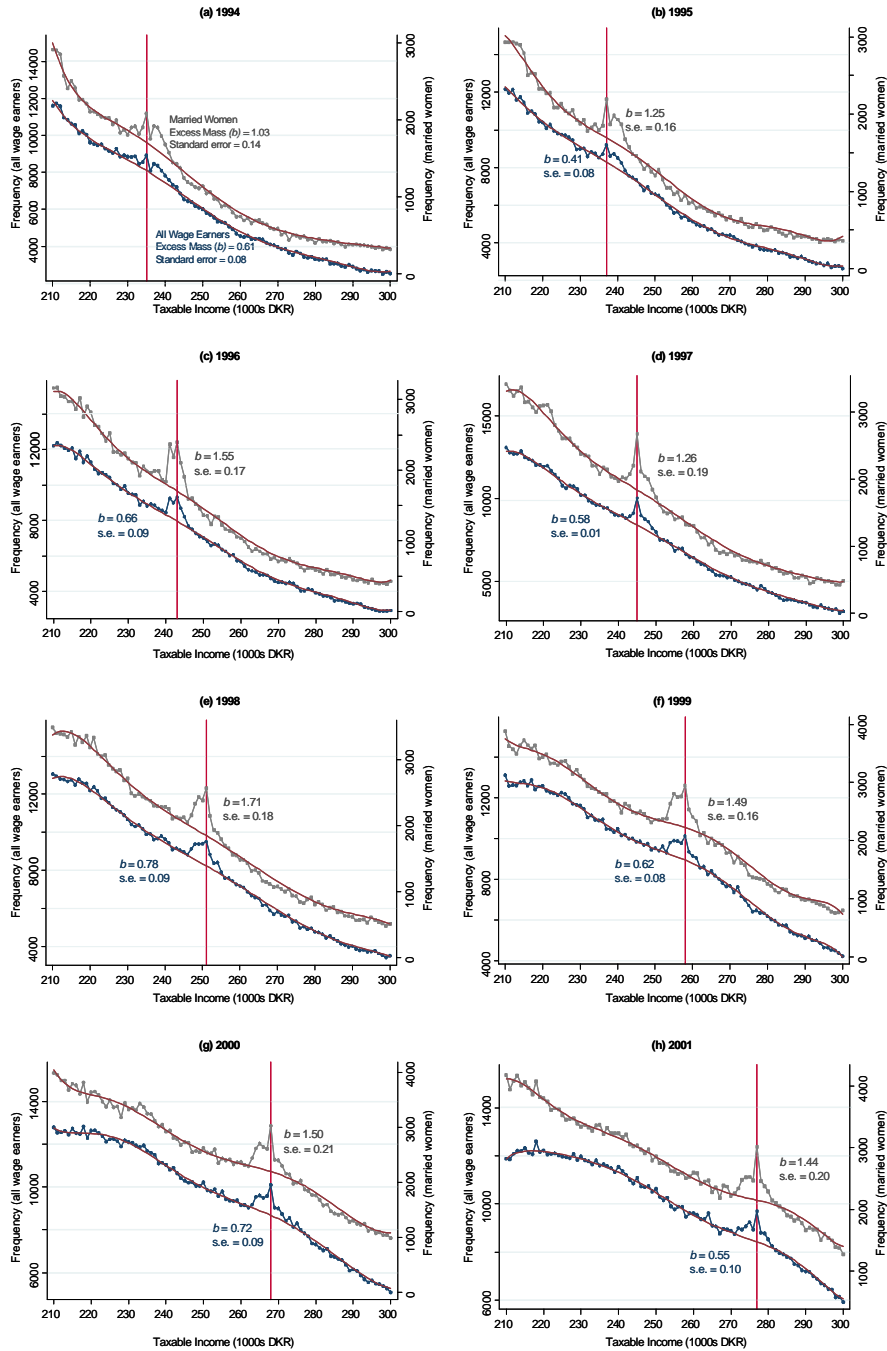
Notes: This figure shows the taxable income distribution around the top tax bracket cutoff (demarcated by the vertical red line at 0) for wage earners between 1994-2001. The series with dots plots a histogram of taxable income (as defined for the top tax base), relative to the top tax cutoff in the relevant year. Each point shows the number of observations in a DKr 1,000 bin. The solid line beneath the empirical distribution is a seventh-degree polynomial fitted to the empirical distribution excluding the points DKr 7,500 or fewer from the cutoff, as in equation (15). The shaded region is the estimated excess mass at the top bracket cutoff, which is 81% of the average height of the counterfactual distribution beneath.

FIGURE 4
Heterogeneity in Bunching at the Top Tax Cutoff



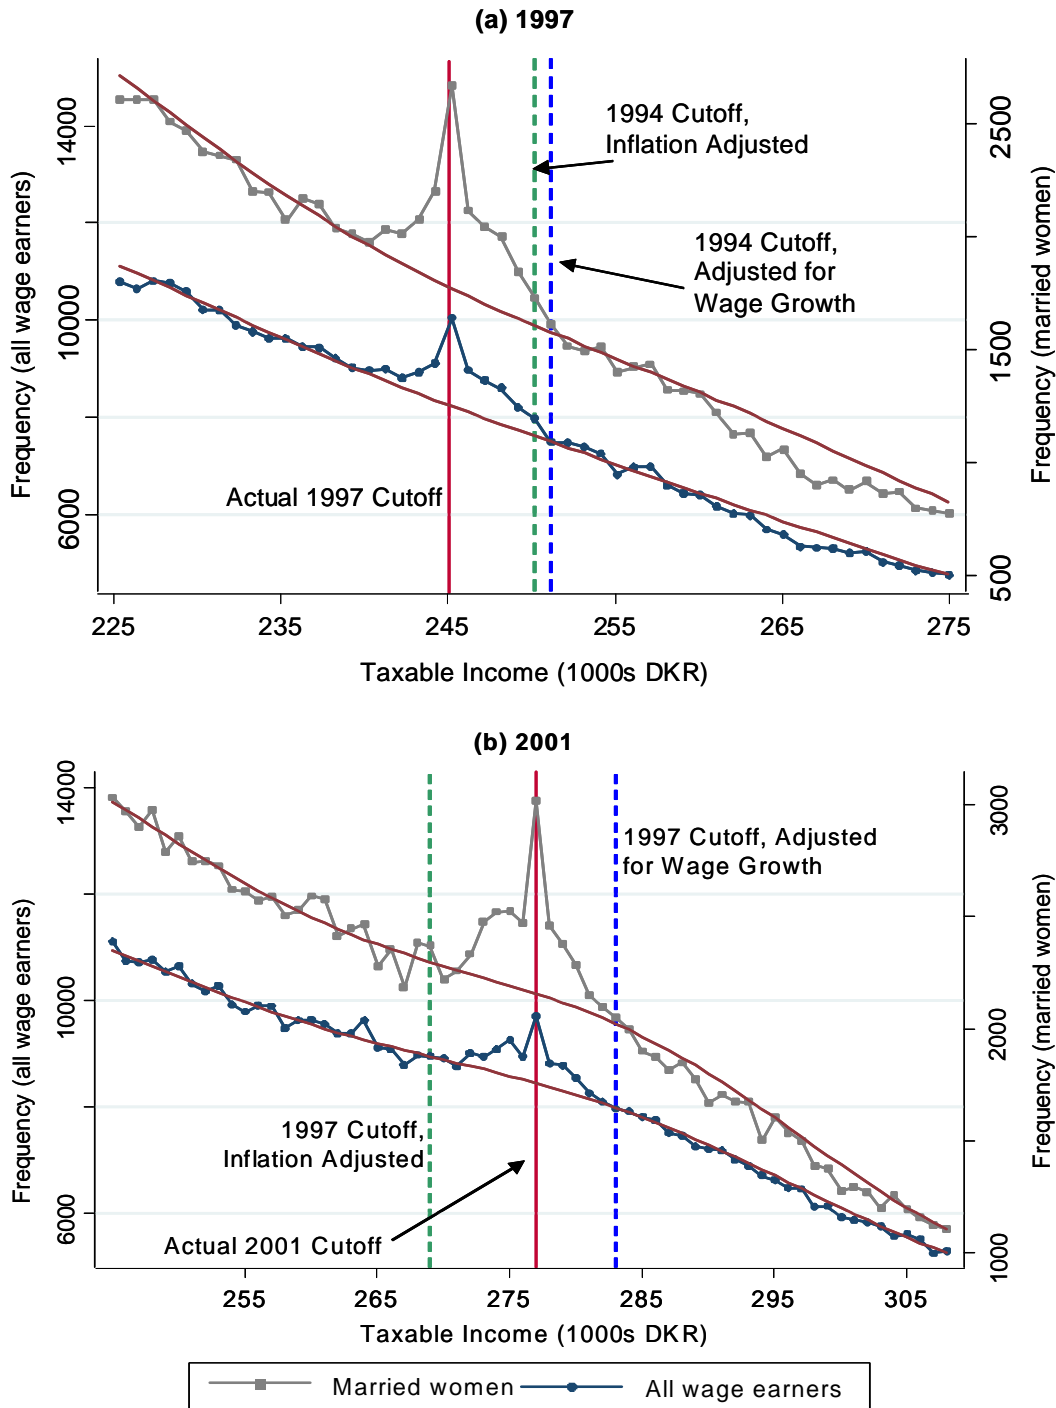
Notes: These figures plot the empirical distributions of taxable income, replicating Figure 3, for four subgroups of the population. Figure (a) considers married women and single men. Figure (b) considers school teachers (ISCO 2331) and the military (ISCO 1013).

FIGURE 5
Income Distributions Around the Top Tax Cutoff, 1994-2001



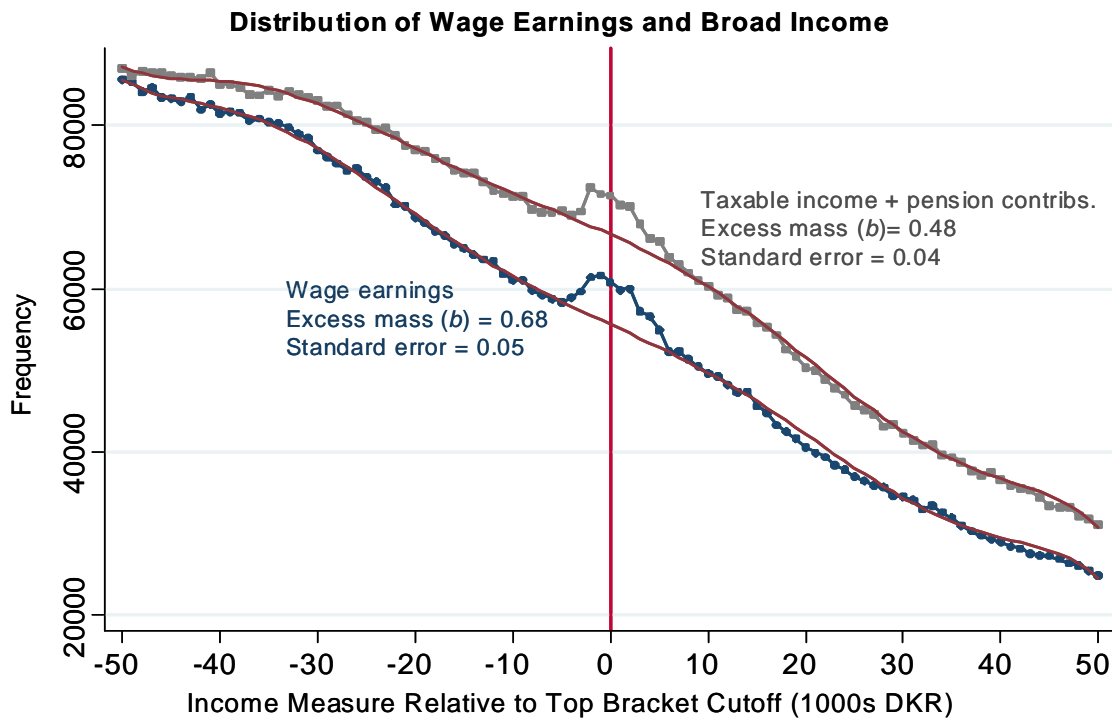
Notes: These figures plot the empirical distribution of taxable income for wage earners and married female wage earners in each year from 1994–2001. In all panels, the upper distribution is for married women and the lower distribution is for all wage earners. The solid vertical lines mark the top tax bracket cutoff (in nominal DKr) in each year. The figure also shows the counterfactual distributions and excess masses, computed as in Figure 3.

FIGURE 6
Distinguishing Changes in Tax Incentives from Inflation and Wage Growth



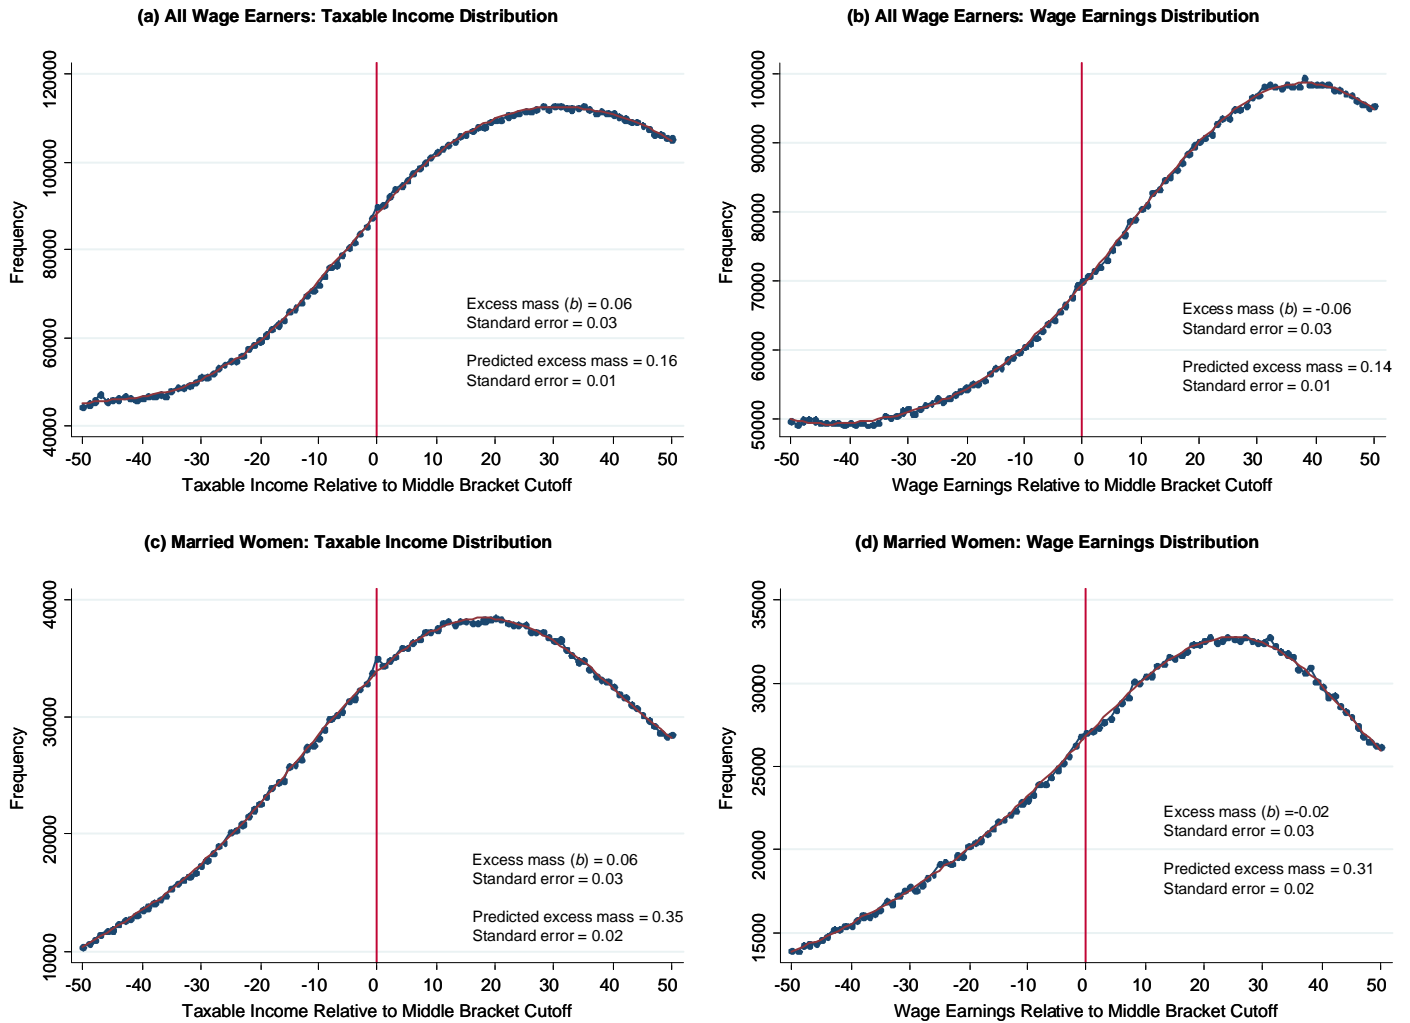
Notes: Panel (a) replicates the income distribution in Figure 5d, zooming in around the top tax bracket cutoff. The location of the bracket cutoff in 1997 is marked with the solid line. The dashed green line shows the level of the 1994 top bracket cutoff adjusted for inflation. The dashed blue line shows the 1994 bracket adjusted for average wage growth. Panel (b) replicates (a) for the year 2001: the dashed vertical lines reflect the 1997 bracket cutoff adjusted for inflation and average wage growth.

FIGURE 7
Distributions of Wage Earnings and Broad Income



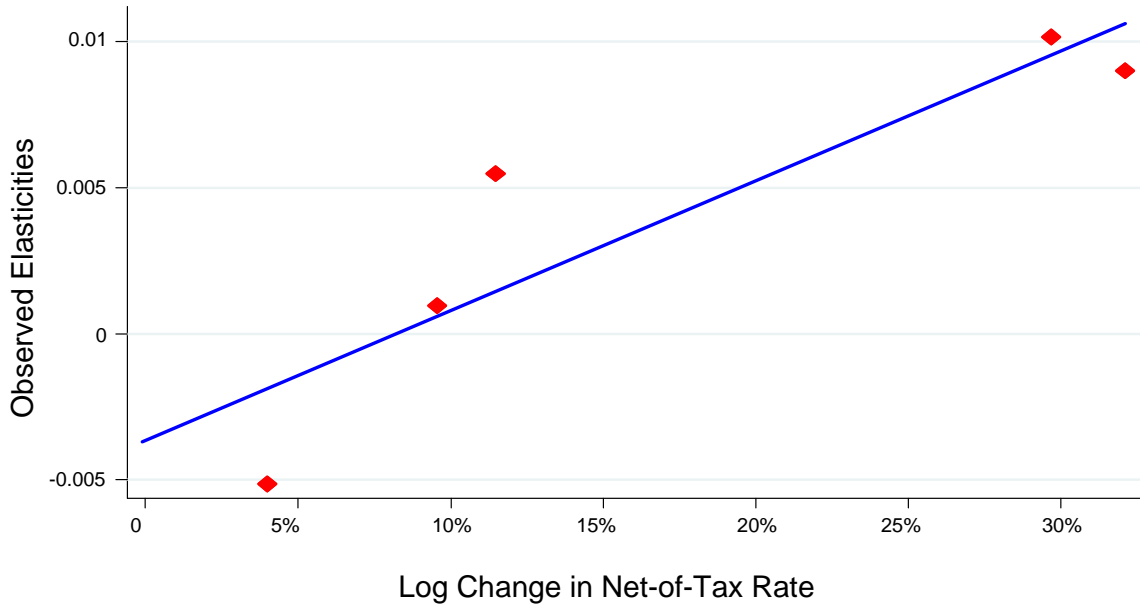
Notes: These figures plot the empirical distribution of wage earnings and broad income around the statutory top tax cutoff (which applies to individuals with 0 net deductions) for the population of wage earners. Broad income is defined as taxable income plus contributions to tax-deductible pension accounts. The figure also shows the counterfactual distributions and excess masses, computed as in Figure 3.

FIGURE 8
Bunching at the Middle Tax Cutoff



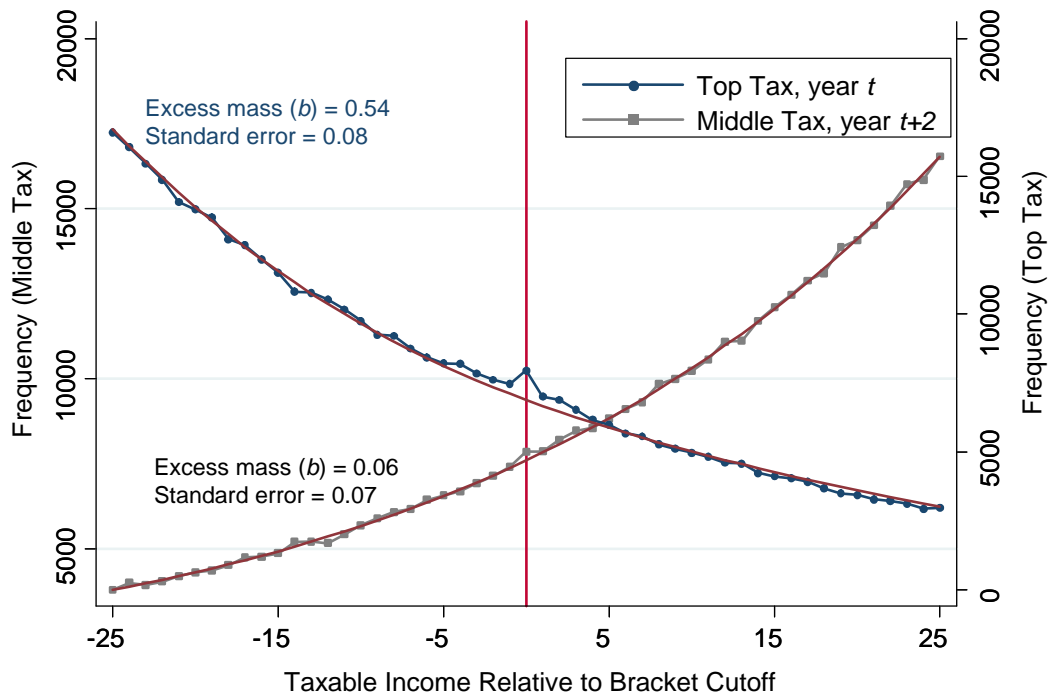
Notes: These figures plot the empirical distributions of taxable income and wage earnings around the middle tax cutoff, where net-of-tax wage rates fall by 11% on average. The figures also show counterfactual distributions and excess masses, computed as in Figure 3. Panel (a) plots the distribution of taxable income (as defined for the middle tax base). Panel (b) plots the distribution of wage earnings, the same variable used in Figure 7. Panels (c) and (d) replicate (a) and (b) for the subgroup of married female wage earners. In each panel, we also report the amount of bunching predicted if the elasticity were the same as that estimated from the amount of bunching at the top bracket cutoff for the corresponding income measure and subgroup.

FIGURE 9
Observed Elasticities vs. Size of Tax Changes



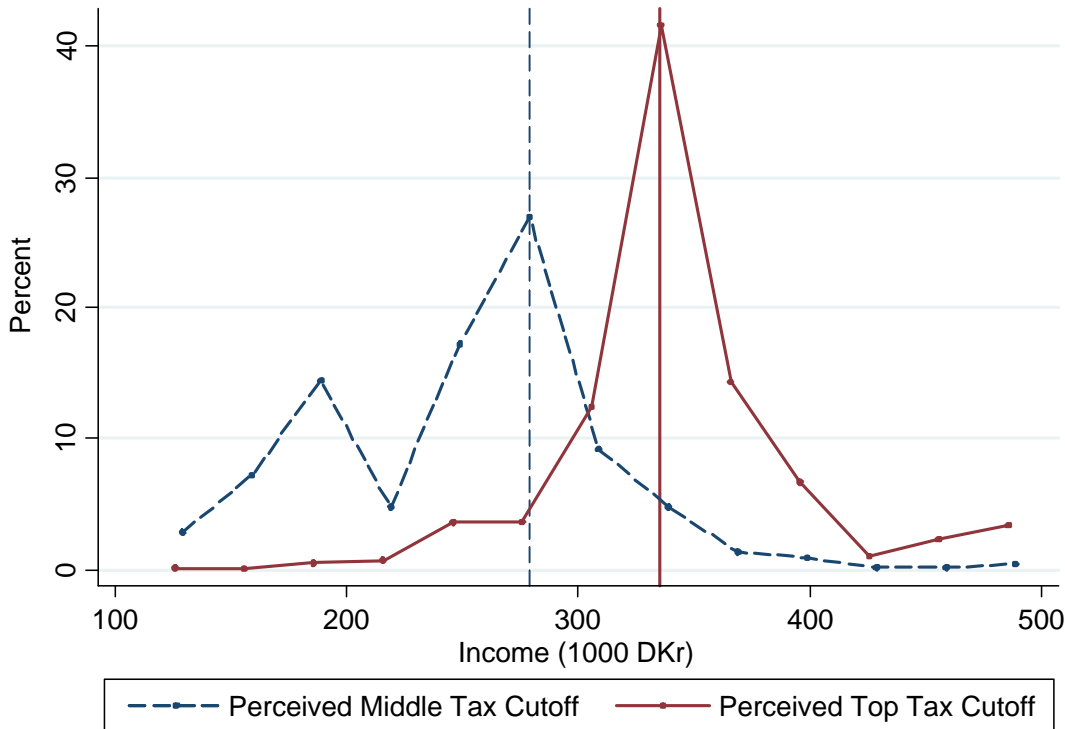
Notes: This figure plots observed elasticities for all wage earners vs. the percent (log) change in the net-of-tax rate used to estimate that elasticity. The blue best-fit line is estimated using OLS. The five points correspond to estimates from small tax reforms (Column 1 of Table 2) and bunching at the middle tax cutoff from 1994-1996 ($\Delta \log(1 - \tau) = 9.5\%$), the middle tax cutoff from 1997-2001 ($\Delta \log(1 - \tau) = 11.6\%$), the top tax cutoff in 1994, 1997, and 1998 ($\Delta \log(1 - \tau) = 29.7\%$), and the top tax cutoff in 1995-1996 and 1999-2001 ($\Delta \log(1 - \tau) = 32.1\%$). The elasticities corresponding to bunching at kinks are calculated using estimates of b as reported in Figure 3 and the formula for the frictionless model in (6).

FIGURE 10
 Bunching for Individuals who Switch Between Top and Middle Kinks



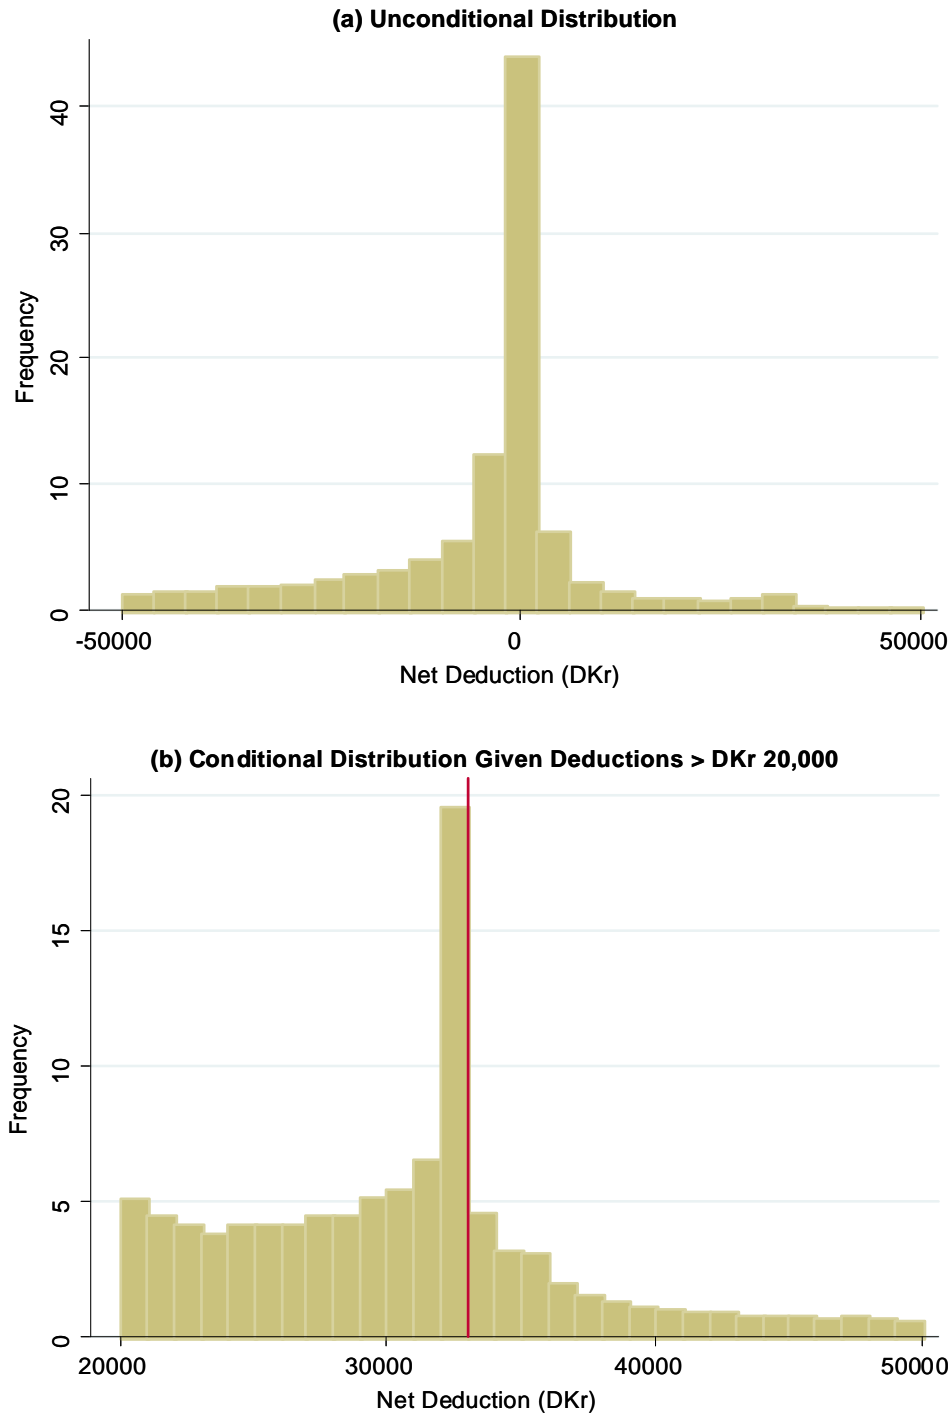
Notes: This figure restricts attention to wage earners who earned within DKr 50,000 of the top tax bracket cutoff in a given year t and within DKr 50,000 of the middle tax bracket cutoff in year $t + 2$. For this fixed group of individuals, we plot the empirical distribution of taxable income in year t around the top bracket cutoff and the distribution of taxable income around the middle tax cutoff in year $t + 2$. The figure also shows the counterfactual distributions and excess masses, computed as in Figure 3.

FIGURE 11
 Survey Evidence: Knowledge about Middle and Top Tax Cutoffs



Notes: This figure plots the distribution of perceived middle and top tax cutoffs from an internet survey of 3,299 members of a union representing public and financial sector employees. Individuals were asked to report the income levels at which they would have to begin paying the middle and top taxes in the 2008 Danish tax code. The figure shows a histogram of the responses for the top tax (solid red line) and middle tax (dashed blue line) cutoffs using bins of DKr 30,000 in width. The bins are centered on the true cutoffs, so that the mode of each distribution represents the fraction of people whose perception of the tax bracket cutoff was within DKr 15,000 of the correct value..

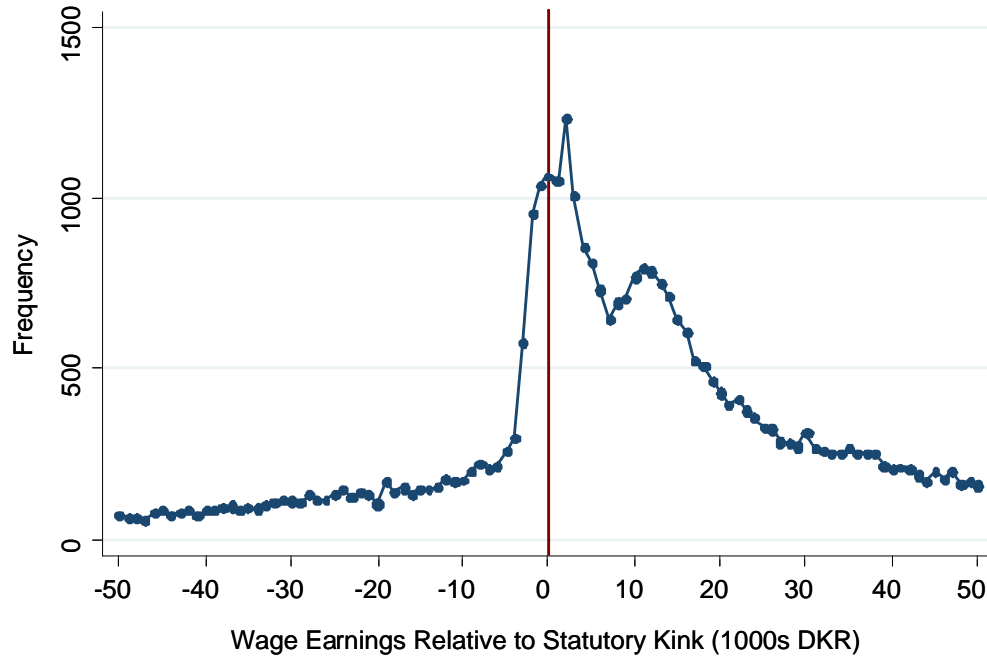
FIGURE 12
Distribution of Net Deductions



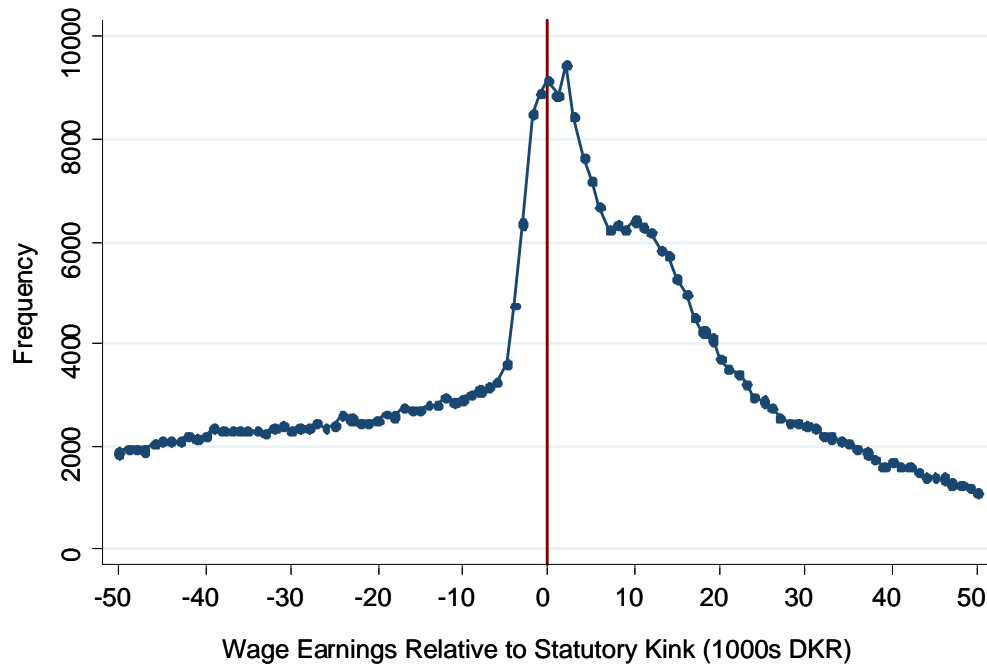
Notes: Panel (a) plots a histogram of net deductions, defined as deductions minus non-wage income relevant for the top tax base. Figure (b) plots a histogram of net deductions between DKr 20,000 and DKr 50,000. To identify bunching in deductions at the pension kink, in Panel B we recenter deductions in each year so that the pension contribution limit in that year equals the average pension contribution limit across the years (DKr 33,000).

FIGURE 13
Teachers' Wage Earnings Distributions

(a) All Teachers

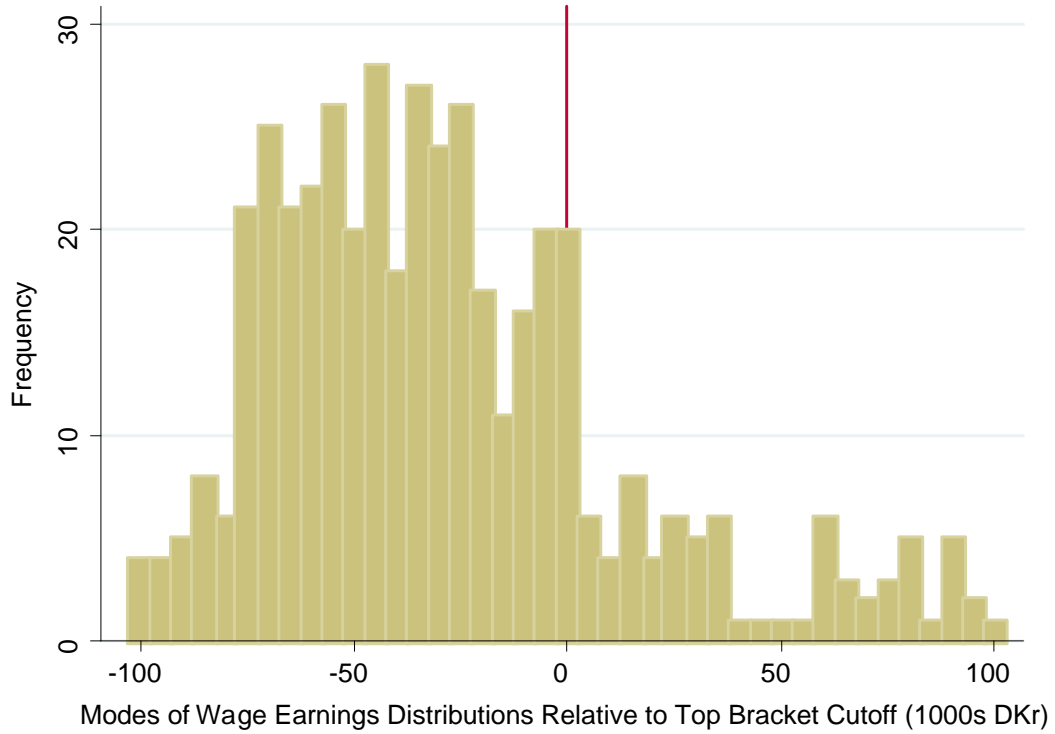


(b) Teachers with Deductions > DKr 20,000



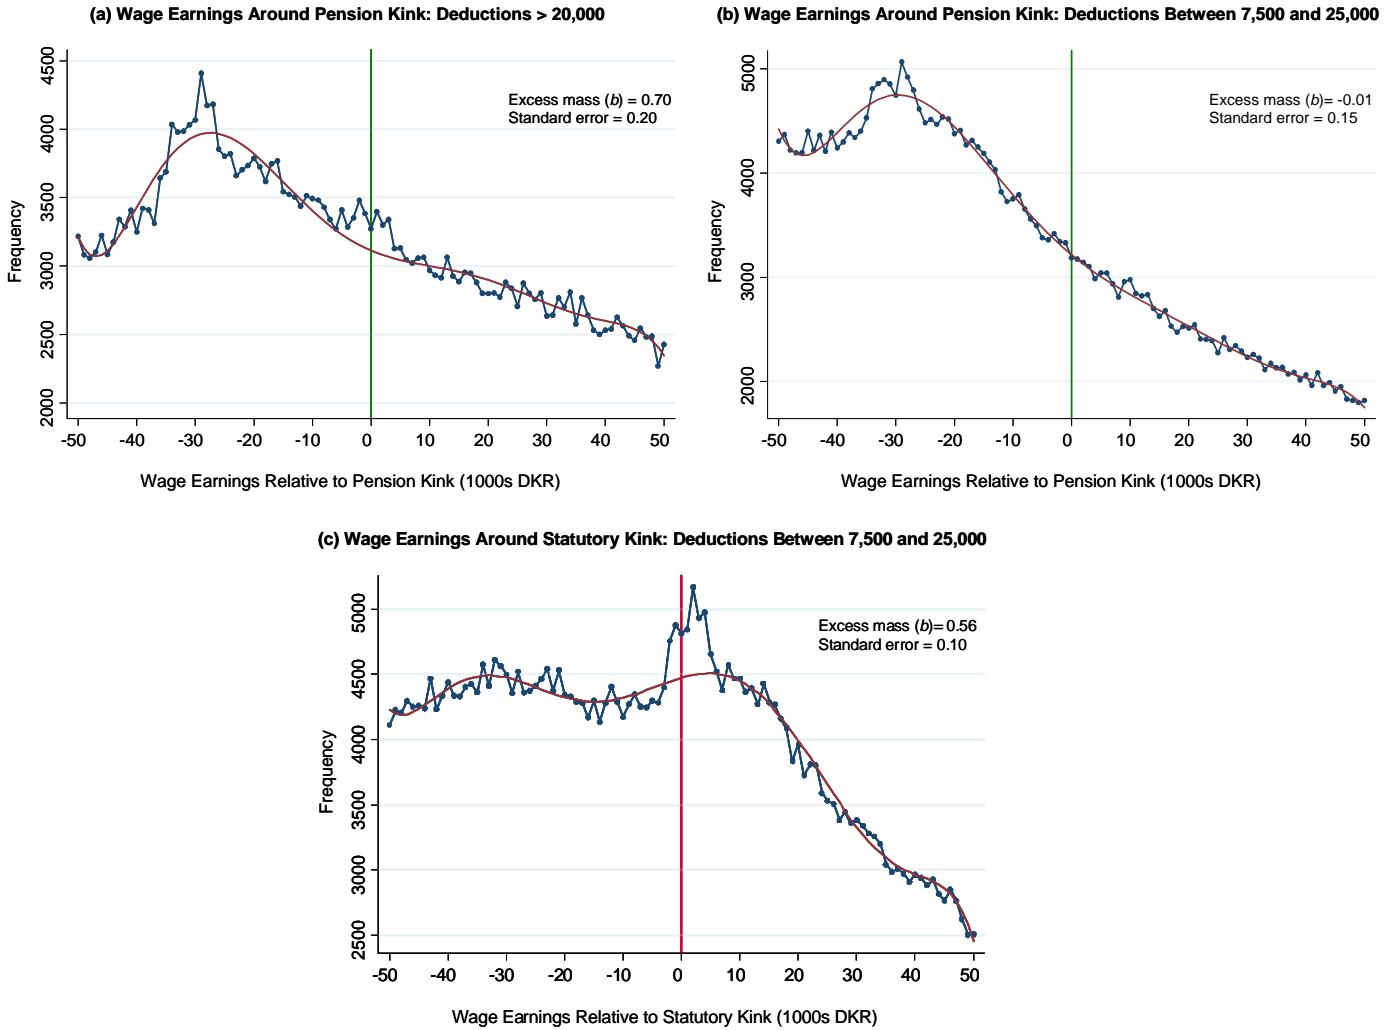
Notes: These two figures plot the empirical distribution of wage earnings around the statutory top tax cutoff in 1994-2001 for (a) all teachers (ISCO 2331) and (b) teachers with net deductions greater than DKr 20,000.

FIGURE 14
Modes of Occupation-Level Wage Earnings Distributions



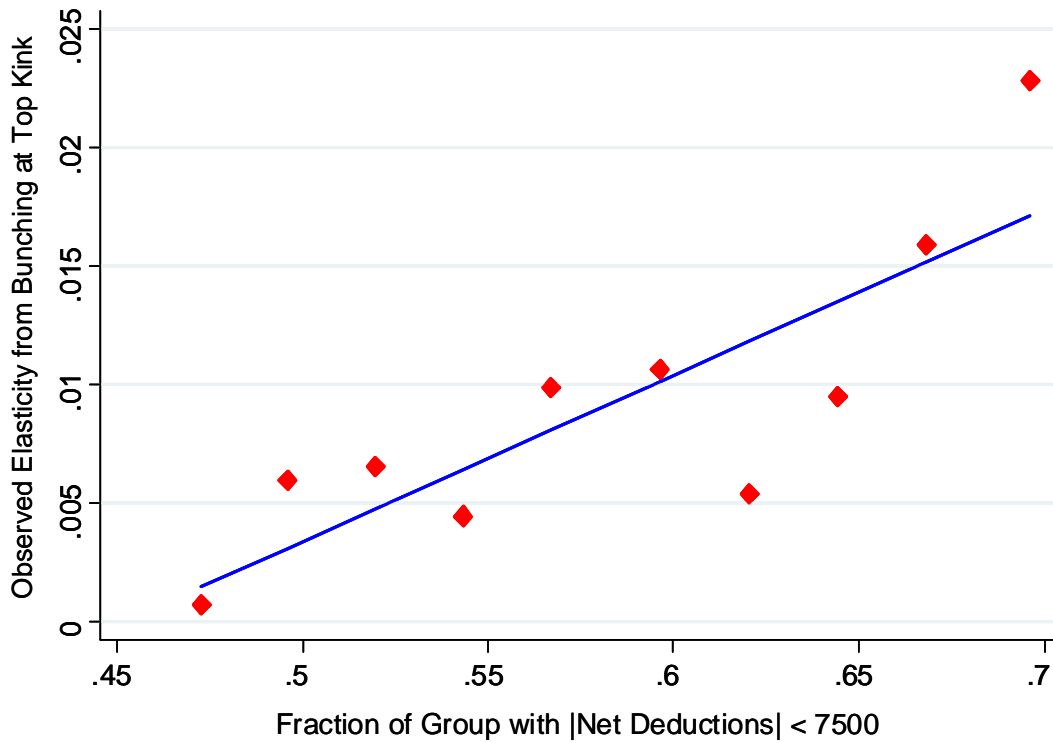
Notes: To construct this figure, we calculate the mode of the wage earnings distribution in each occupation-year cell, defined as the DKr 5,000 bin with the most individuals in that occupation-year. Occupations are defined by 4 digit ISCO codes. The figure shows a histogram of these modes, excluding occupations with fewer than 7000 workers.

FIGURE 15
Individual vs. Firm Bunching at the Pension Kink



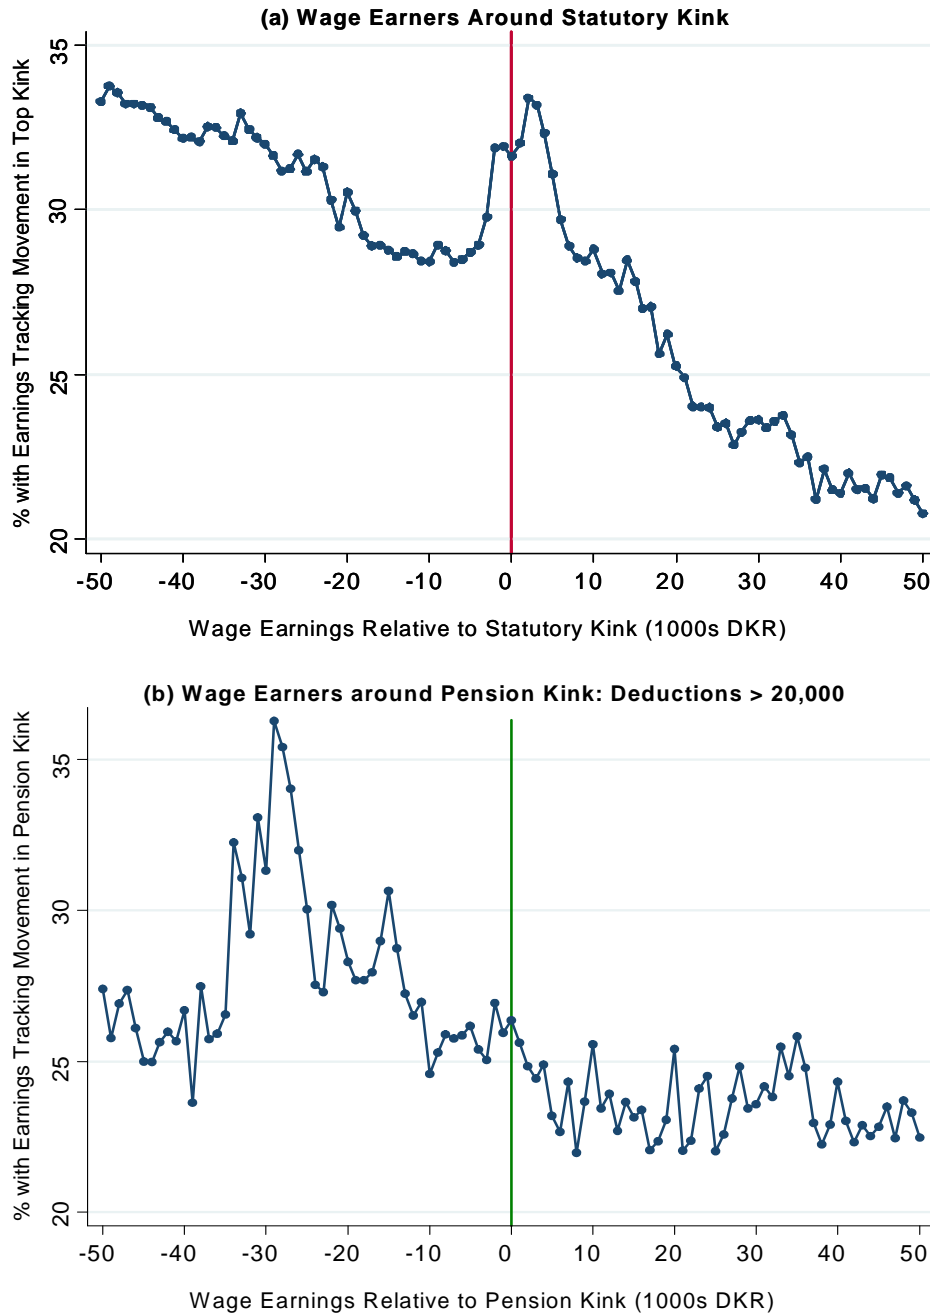
Notes: Panel (a) plots the distribution of wage earnings relative to the pension kink (demarcated by the green vertical line) for wage earners with greater than DKr 20,000 of net deductions. The pension kink is defined as the top tax bracket cutoff plus the maximum tax-deductible pension contribution in each year. Panel (b) replicates (a) for wage earners with between DKr 7,500 and DKr 25,000 of net deductions. Panel (c) plots the distribution of wage earnings relative to the statutory top kink (demarcated by the red vertical line) for wage earners with between DKr 7,500 and DKr 25,000 in net deductions. The figure also shows the counterfactual distributions and excess masses, computed as in Figure 3.

FIGURE 16
Observed Elasticities vs. Scope of Tax Changes



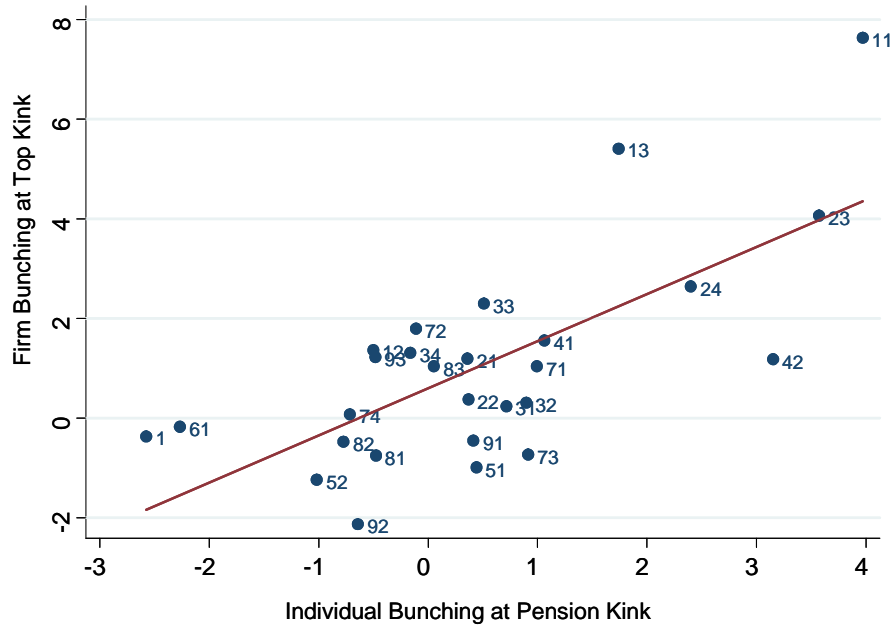
Notes: To construct this figure, we first calculate the fraction of individuals with net deductions less than DKr 7,500 in magnitude in each age-gender-marital status-year cell. We then group individuals into 10 equal-width bins based on the fraction with small deductions in their group as described in the text. We estimate the excess mass at the top kink as in Figure 3 and apply equation (6) to calculate observed elasticities for each of the ten groups. The figure shows a scatter plot of the observed elasticities vs. the fraction with small deductions in the 10 bins. The blue best-fit line is estimated using OLS.

FIGURE 17
Dynamics of Earnings Around the Top Tax Cutoff



Notes: These figures show how the propensity to track the movement in the top tax cutoff across years varies across individuals. To construct Panel (a), we first divide individuals into bins of DKr 1000 in wage earnings in a given year t , and calculate the fraction in each bin whose change in wage earnings from a year t to $t + 2$ falls within DKr 7,500 of the movement in the top tax bracket cutoff from year t to $t + 2$. Panel A plots this fraction for wage earnings bins around the statutory top tax cutoff. Panel (b) replicates (a) for the pension kink, restricting the sample to wage earners with net deductions greater than DKr 20,000.. It shows the fraction of individuals whose change in wage earnings falls within DKr 7,500 of the movement in the pension kink for wage earnings bins around the pension kink.

FIGURE 18
Correlation Between Individual and Firm Bunching

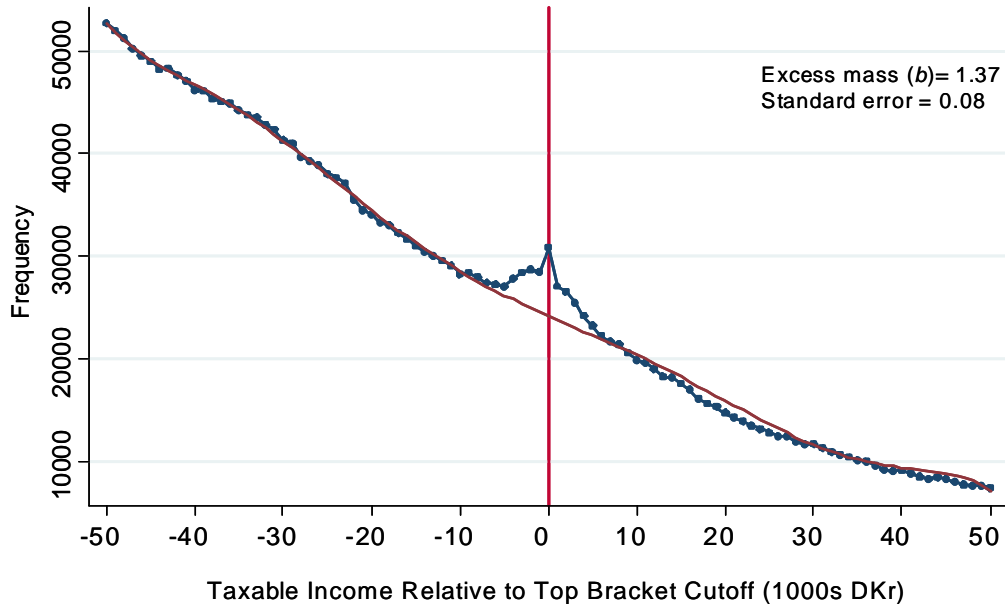


ISCO Code	Occupation Description	No. of Workers
1	Military	24,557
11	Legislators and senior officials	5,918
12	Corporate managers	54,154
13	General managers	3,238
21	Physical, mathematical and engineering science professionals	42,645
22	Life science and health professionals	30,058
23	Teaching professionals	106,735
24	Other professionals	72,240
31	Physical and engineering science associate professionals	68,803
32	Life science and health associate professionals	73,960
33	Teaching associate professionals	77,032
34	Other associate professionals	147,951
41	Office clerks	237,735
42	Customer service clerks	32,853
51	Personal and protective service workers	228,658
52	Models, sales persons, and demonstrators	74,005
61	Skilled agricultural and fishery workers	14,198
71	Exaction and related trades workers	95,359
72	Metal, machinery and related trades workers	110,937
73	Precision, handicraft, printing and related trades workers	11,786
74	Other craft and related trades workers	20,086
81	Stationary plant and related operators	10,968
82	Machine operators and assemblers	106,556
83	Drivers and mobile plant operators	36,318
91	Sales and services elementary occupations	99,551
92	Agricultural, fishery and related labourers	9,431
93	Mining, construction, manufacturing, and transport	75,516

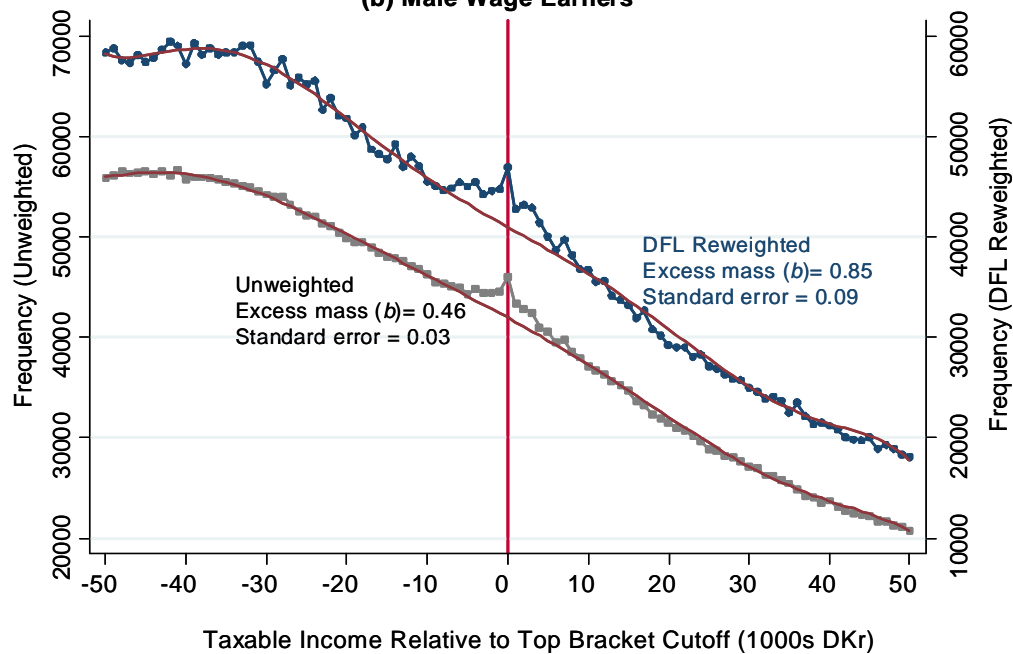
Notes: This figure plots the amount of firm bunching vs. the amount of individual bunching for all International Standard Classification of Occupation codes at the two digit level. Both firm and individual bunching are estimated on the subgroup of individuals with net deductions greater than DKr 20,000, as in Figure 15a. Individual bunching is the excess mass at the pension kink for this group, while firm bunching is the excess mass at the statutory top tax cutoff for the same group.

FIGURE 19
Male vs. Female Wage Earners: Effects of Occupational Heterogeneity

(a) Female Wage Earners

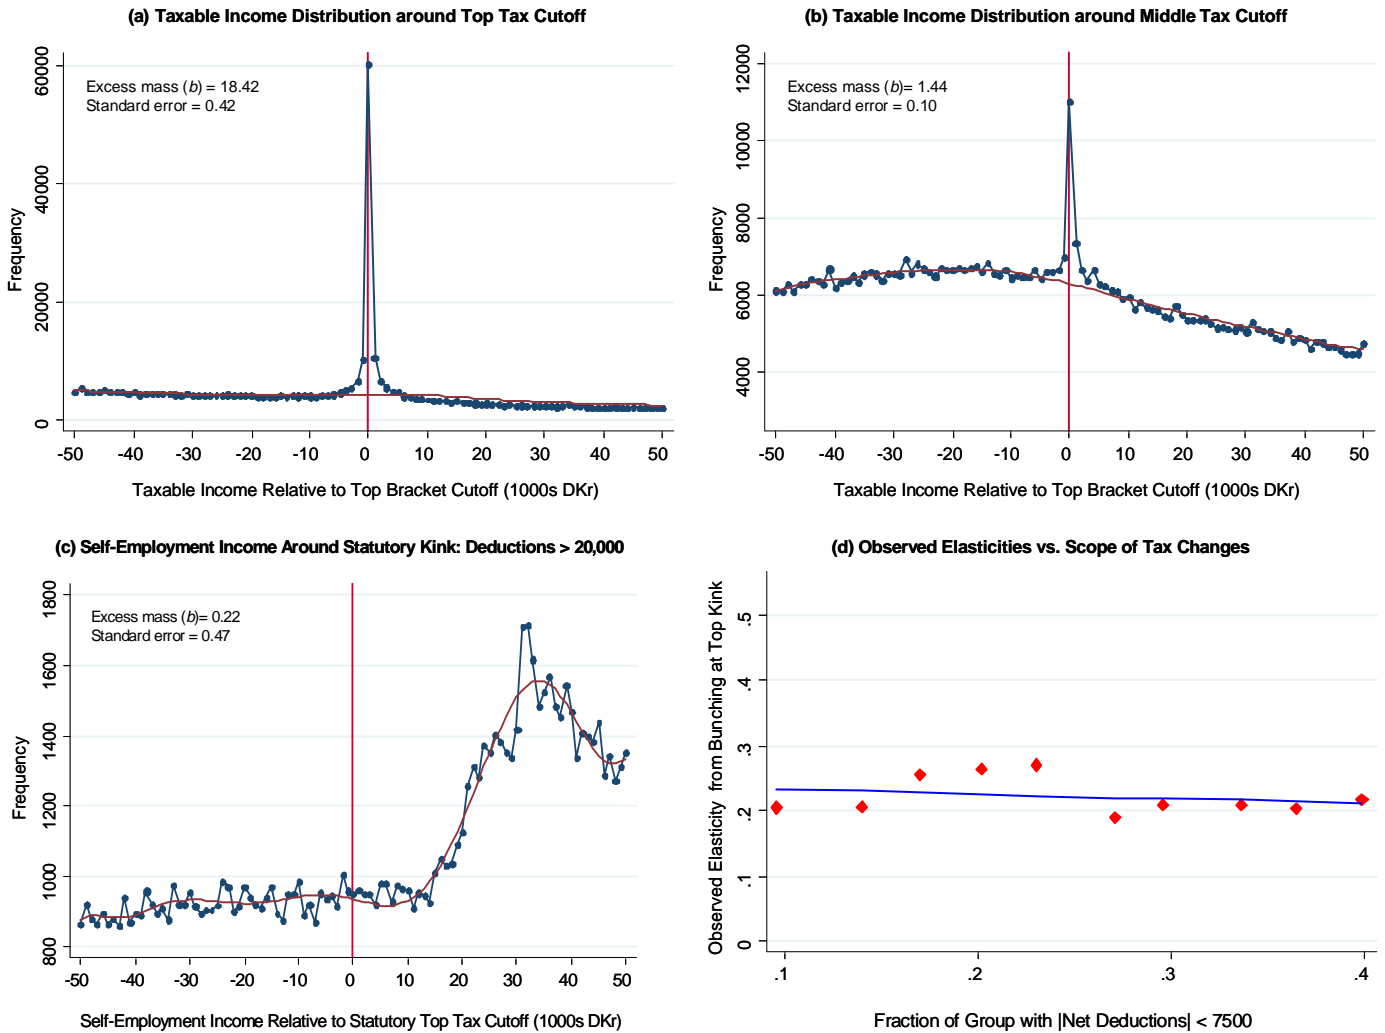


(b) Male Wage Earners



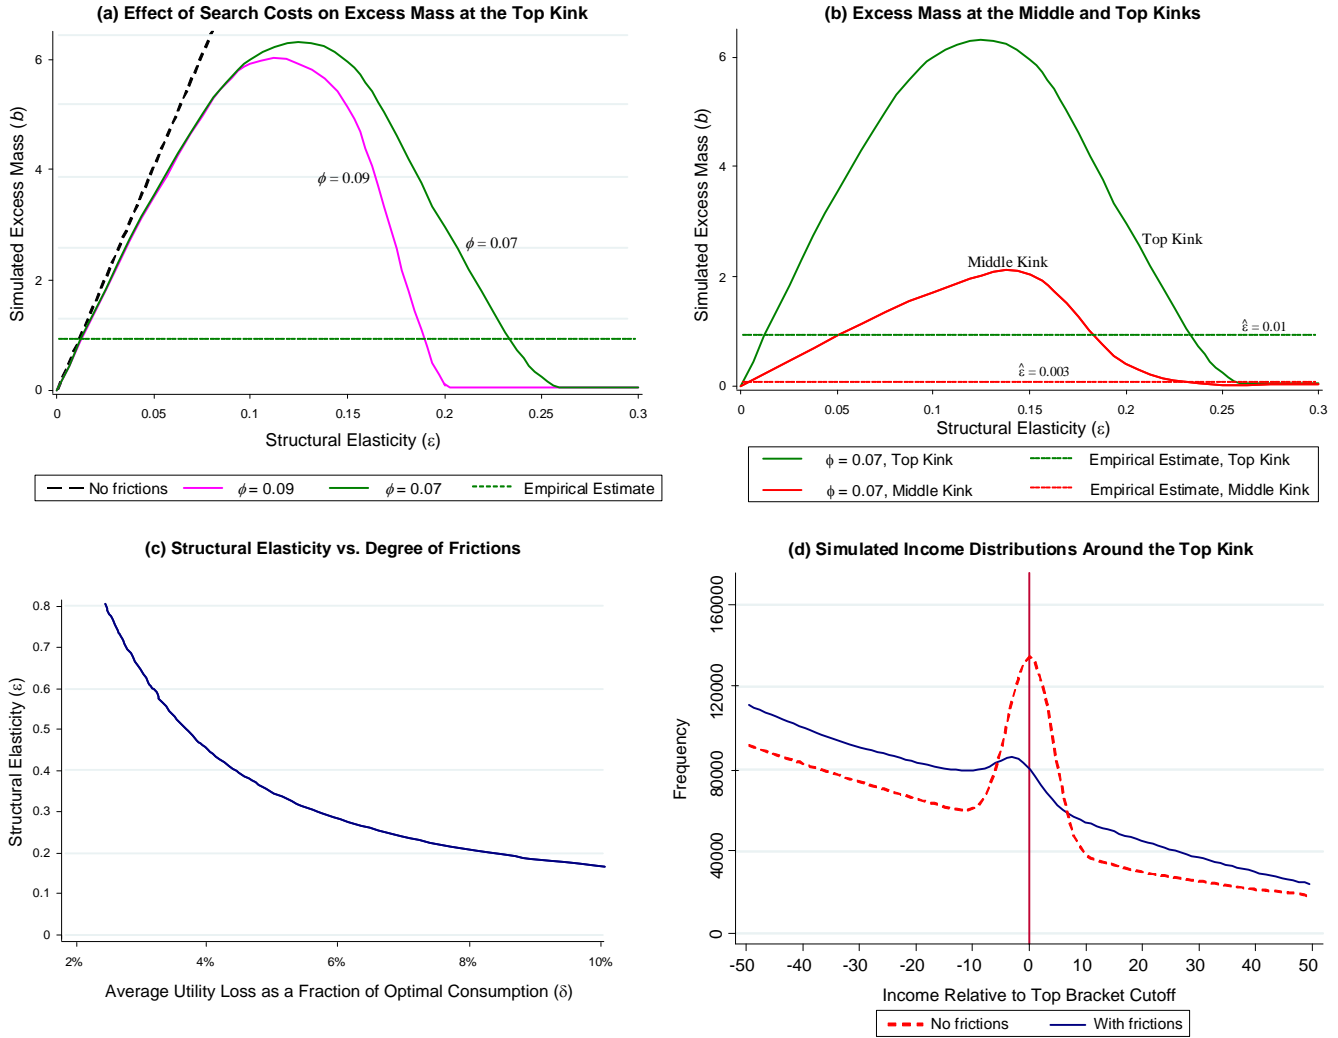
Notes: These figures plot the empirical distributions of taxable income around the top tax cutoff for (a) female wage earners and (b) male wage earners. The series in grey squares in Panel B shows the raw distribution of taxable income for men. The series in blue circles shows reweights the observations for men to match the occupational distribution of women (defined by 4 digit ISCO codes). Following DiNardo, Fortin, and Lemieux (1996), we reweight an observation in occupation i by $\frac{p_i}{1-p_i}$, where p_i is the probability that a wage earner in occupation i is female. The figure also shows the counterfactual distributions and excess masses, computed as in Figure 3.

FIGURE 20
Self-Employed Individuals



Notes: These figures include only individuals who report positive self-employment income. Panels (a) and (b) plot the taxable income distribution around the top and middle cutoffs from 1994-2001. Panel (c) plots the distribution of realized self-employment income around the statutory top tax cutoff for individuals with net deductions greater than 20,000. Panel (d) replicates Figure 16 for individuals with positive self-employment income, with the y axis scaled to have the same range relative to the mean observed elasticity as in Figure 16.

FIGURE 21
Calibration of Model and Bound on Structural Elasticity



Notes: This figure shows how we calibrate the model to match the empirically estimates of bunching at the top and middle kinks. Panel (a) plots the simulated relationship between the structural elasticity ε and the observed bunching at the top kink for two values of the search cost: $\phi = 0.07$, and $\phi = 0.09$. Panel (b) plots the simulated amount of bunching at the top and middle kinks vs. ε for the $\phi = 0.07$ case. The horizontal lines show the empirically observed amount of bunching at the two kinks. Panel (c) shows the value of ε that fits the data for each value of δ , the average utility loss from failing to choose optimal hours measured as a fraction of optimal consumption. Panel (d) displays the simulated equilibrium taxable income distributions without frictions ($\phi = 0$) and with frictions ($\phi = 0.049$, $\gamma = 1.015$) when $\varepsilon = 0.34$. See Appendix C for details of the numerical methods used to generate these figures.