

NBER WORKING PAPER SERIES

MICRO VS MACRO LABOR SUPPLY ELASTICITIES:
THE ROLE OF DYNAMIC RETURNS TO EFFORT

Henrik Kleven
Claus Thustrup Kreiner
Kristian Larsen
Jakob Egholt Sørensen

Working Paper 31549
<http://www.nber.org/papers/w31549>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2023, Revised April 2024

We thank Raj Chetty, Jeppe Druedahl, John Friedman, Peter Ganong, Nathan Hendren, Pat Kline, Wojciech Kopczuk, Ilyana Kuziemko, Camille Landais, Felipe Lobel, Benjamin Lockwood, Emi Nakamura, Richard Rogerson, Benjamin Schoefer, Johannes Spinnewijn, Damián Vergara, Danny Yagan, and Owen Zidar for helpful comments and discussions. The activities of CEBI are financed by the Danish National Research Foundation, grant DNRF134. The views expressed herein are those of the authors 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Henrik Kleven, Claus Thustrup Kreiner, Kristian Larsen, and Jakob Egholt Sørensen. 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.

Micro vs Macro Labor Supply Elasticities: The Role of Dynamic Returns to Effort
Henrik Kleven, Claus Thustrup Kreiner, Kristian Larsen, and Jakob Egholt Sørensen
NBER Working Paper No. 31549
August 2023, Revised April 2024
JEL No. C1,D6,E6,H2,H3,J2,J3

ABSTRACT

A key contention in economics is the discrepancy between micro and macro elasticities of labor supply with respect to marginal tax rates. We revisit this question, focusing on the role of dynamic returns to effort among top earners. We develop a new model of earnings responses to taxes in the presence of dynamic returns. In this model, the returns to effort are delayed and mediated by job switches such as promotions within firms or movements between firms. Short-run micro elasticities are attenuated relative to the true long-run macro elasticity. We proceed by providing two main empirical analyses using rich administrative data from Denmark. The first part presents descriptive evidence on earnings and hours-worked patterns over the lifecycle that confirm the predictions of the theoretical model. The second part presents quasi-experimental evidence on earnings responses to taxes using discrete job switches. The empirical strategy is informed by the theoretical model, according to which job switches can be used to (partially) identify the macro elasticity of labor supply. The evidence shows that, at the top of the distribution, macro elasticities are much larger than micro elasticities due to dynamic compensation effects.

Henrik Kleven
Department of Economics
Princeton University
238 Julis Romo Rabinowitz Building
Princeton, NJ 08544
and CEPR
and also NBER
kleven@princeton.edu

Claus Thustrup Kreiner
Department of Economics
University of Copenhagen and CEBI
Øster Farimagsgade 5
Building 35
DK-1353 Copenhagen K
Denmark
ctk@econ.ku.dk

Kristian Larsen
University of Copenhagen and CEBI
Department of Economics
Øster Farimagsgade 5
Building 35
DK-1353
Copenhagen K
Denmark
kuol@econ.ku.dk

Jakob Egholt Sørensen
University of Copenhagen and CEBI
Øster Farimagsgade 5
Building 26
DK-1353 Copenhagen K
Denmark
jes@econ.ku.dk

1 Introduction

The impact of taxes on labor supply and earnings is critical for assessing the equality-efficiency trade-off and optimal redistribution. The strength of these responses is typically measured by the elasticity of earnings with respect to taxes. Importantly, the welfare-relevant elasticity captures long-run responses, accounting for tax-induced changes in the full lifetime trajectory of earnings. Such responses are challenging to estimate and there is no consensus on plausible magnitudes. Even if much microeconomic evidence points to small elasticities, many economists contend that the true long-run effect of taxes is large due to dynamic effects (e.g., [Prescott 2004](#)).

This debate relates to the stark difference between micro and macro elasticities of labor supply. Micro estimates — typically using tax reforms as quasi-experiments — tend to be small. Macro estimates on the other hand — typically based on structural estimations or calibrations — tend to be large. To illustrate the extent of disagreement between these two research strands, consider the Laffer rate on top earners. The top-income Laffer rate in the US is close to 80% if the elasticity is 0.2 (a typical micro estimate), but only 40% if the elasticity is 1 (a typical macro estimate).¹ This range is too large to provide useful policy guidance.

Which approach is right and which is wrong? Our starting point is that both are right and both are wrong. Micro studies are based on research designs that allow for causal identification, but the approach only captures short-run effects and may miss important dynamic mechanisms. Macro studies are model-dependent and may be associated with specification bias, but they allow for potentially relevant dynamic responses. The goal of this paper is to develop a quasi-experimental approach that is better able to capture welfare-relevant, long-run elasticities.

We are particularly interested in the elasticity at the top of the income distribution, among salaried career workers. Such top earners represent a large fraction of income and tax revenue, making them critical to tax design. A key challenge to estimating their responsiveness to taxes is the presence of dynamic returns to effort. Consider an example close to home: top academics. An academic career consists of doing research and building a publication record, which may eventu-

¹These numbers are based on the Laffer rate formula $\tau = 1 / (1 + \varepsilon a)$, where ε is the earnings elasticity and a is the Pareto parameter ([Diamond 1998](#); [Saez 2001](#)). The Pareto parameter is about 1.5 in the US. An elasticity of 1 is, if anything, a conservative assessment of the macro literature (see e.g., [Prescott 2004](#); [Rogerson and Wallenius 2009](#); [Keane and Rogerson 2015](#)). For example, the cross-country calibration study by [Prescott \(2004\)](#) argues for a (Hicksean) hours-worked elasticity much larger than 1.

ally lead to promotions within a department or outside offers from other departments. The link between effort and earnings is delayed and discrete, centered around promotions or firm switches. We posit that such dynamic and discrete returns characterize most top professions. Standard quasi-experimental research designs are largely uninformative in the presence of dynamic returns. They rely on models where outcomes respond almost immediately to incentives such as in the static and frictionless model of hourly-paid workers.

Our agenda relates to an idea previously studied in macro and structural labor economics: the effect of effort on human capital accumulation via learning-by-doing and on-the-job training. Structural estimations of human capital models are consistent with large long-run elasticities (see e.g., [Keane 2011](#); [Keane and Rogerson 2015](#)).² While human capital accumulation is one way of generating dynamic returns, a variety of other mechanisms may be at play. In the example of top academics, it is not *a priori* clear if wages increase over the career path due to changes in productivity or because discrete performance evaluations reward historical output. We take the latter view, the implication of which is that the lifecycle profile of earnings is a step function with discrete changes at job switches such as occupation or firm switches. This idea is related to a literature in labor economics showing that job-to-job mobility is central for earnings growth over the lifecycle either through gains in the job-match component of wages (e.g., [Topel and Ward 1992](#); [Pavan 2011](#)) or through mobility to firms with higher wage premia (e.g., [Abowd, Kramarz, and Margolis 1999](#); [Card, Heining, and Kline 2013](#); [Card, Cardoso, and Kline 2016](#); [Card, Cardoso, Heining, and Kline 2018](#)).³ We develop a new model that highlights the role of dynamic returns realized at the point of switching and investigate the implications for the estimation of behavioral responses to taxes.

Our model distinguishes between *realized* earnings and *latent* earnings (effort). Workers make effort choices based on their productivity and taxes, taking into account that higher effort generates higher earnings with a delay. Realized earnings change only at discrete job events — such as switches between occupations or firms — at which time realized and latent earnings are realigned. We consider a benchmark model where the probability of switching is exogenous and an extension where this probability is endogenized. The standard labor supply model obtains as a special case with a switching probability of one, in which case the elasticity of true effort η (which governs the

²[Best and Kleven \(2013\)](#) develop a theory of optimal taxation in a setting where effort affects future wages through human capital accumulation.

³A macro literature also provides evidence on the importance of job switches for wage and earnings growth (e.g., [Grigsby, Hurst, and Yildirmaz 2021](#); [Guvenen, Karahan, Ozkan, and Song 2021](#)).

long-run macro elasticity) is identical to the elasticity of realized earnings ε (the short-run micro elasticity). Outside this limit case, the macro elasticity is larger than the micro elasticity. Allowing for heterogeneity in both structural elasticities η and switching probabilities λ , we characterize the conditions under which the macro elasticity can be point identified or partially identified using responses among short-run switchers. The macro elasticity can be point identified when η and λ are orthogonal and partially identified when η and λ are correlated. The two cases can be separated based on the relationship between observed micro elasticities and the timing of switches following a tax reform.

The empirical part of the paper leverages Danish administrative data to verify the predictions of the model and identify the long-run macro elasticity. The data are employer-employee matched and contain detailed occupation codes, allowing us to observe job switches at a granular level. We start by providing descriptive evidence on earnings and hours-worked patterns over the lifecycle, highlighting four facts of the data. First, earnings are strongly related to past hours worked, conditional on current hours worked. Considering workers at advanced career stages, past hours is a stronger predictor of earnings than current hours. Second, contemporaneous changes in hours worked and earnings are virtually unrelated at the top of the distribution, but not at the bottom. That is, while the hours-earnings relationship at the bottom (such as for cleaners, cashiers, and hotel porters) is consistent with the textbook model of hourly-paid workers, the relationship higher up in the distribution (such as for engineers, lawyers, and managers) is consistent with a model of salaried workers for whom earnings and effort feature little correlation period by period. Third, for workers who reach the top of the distribution, the lifecycle profile of earnings is discrete, driven by jumps at job switches and inaction between switches. In fact, between-job variation accounts for almost 95% of the total variation in earnings over the lifecycle. Finally, based on event studies of promotions — defined as switches to job cells with higher median earnings — we show that individual earnings jump discretely at promotion events while hours worked are smooth.⁴ These four facts are consistent with our theoretical model, but not with standard models.

Informed by our model and descriptive evidence, we provide a quasi-experimental study of earnings elasticities using firm \times occupation switchers. The analysis is based on a recent tax reform experiment in Denmark: a reduction in the marginal tax rate above an income threshold located at around the 70th percentile of the distribution. The tax rate reduction was large, about 11pp.⁵ We

⁴Related, [Bronson and Thoursie \(2021\)](#) show that discrete earnings jumps within firms account for a large fraction of the gender gap in earnings.

⁵Given the tax reform was permanent, our estimates are most naturally interpreted as capturing Hicksian elasticities,

estimate behavioral responses using difference-in-differences and triple-differences approaches. Considering the full population of treated workers, we find clear and precisely estimated earnings responses to taxes. However, the average response is modest in elasticity terms. We then split the data into job movers and job stayers, showing that the small average response masks striking heterogeneity: job movers feature large responses — an elasticity of 0.4-0.5 — while job stayers feature no responses. Because a minority of workers switch in any given year, the large switcher elasticity is consistent with a small average elasticity. We also consider the effect of the tax reform on the probability of switching, finding no effect on this margin. This is consistent with our theoretical model in which the switching probability, even when endogenously set by firms, does not respond to income taxes.

The central thesis of our paper is that earnings responses among short-run switchers can be used to uncover the long-run macro elasticity in the population. As mentioned, point identification requires orthogonality between structural elasticities η and switching probabilities λ , the implication of which is that the observed switcher elasticity ε is constant in the timing of switching. Estimating impacts by the timing of switching jobs, we show that the time profile of the switcher elasticity is virtually flat, consistent with point identification. We provide two additional pieces of evidence that speak to this point. First, we investigate if switcher characteristics respond to the reform — i.e., if any differences between switchers above and below the treatment threshold change from before to after the reform. Looking at a wide range of switcher characteristics, we do not find effects on any of them. As a result, controlling for switcher characteristics in the empirical specification hardly affects the estimates. Second, we restrict the sample to plausibly exogenous switches, namely those triggered by mass layoffs. Mass-layoff switchers feature similar earnings responses to the tax reform as the full sample of switchers.⁶ Taken together, this set of findings suggests that our estimates are not biased by selection into switching.

Our paper contributes to several literatures in public finance, labor, and macroeconomics. First, we contribute to an enormous body of work estimating labor supply elasticities with respect to tax incentives, as reviewed by [Blundell and MaCurdy \(1999\)](#) and [Saez, Slemrod, and Giertz \(2012\)](#).

Summarizing the microeconomic evidence, [Saez, Slemrod, and Giertz \(2012\)](#) argue that “the

not Frisch elasticities. A conceptual challenge to all tax reform studies, however, is that “permanent” tax changes are never truly permanent and behavioral responses may be influenced by (unobserved) expectations about future tax reforms. In this sense, it is conceivable that tax reform studies capture a mix of Hicksian and Frisch elasticities.

⁶Importantly, the positive earnings response to lower taxes for mass-layoff switchers is based on our quasi-experimental design, which compares treated and untreated mass-layoff switchers from before to after the reform. As we show, this earnings response is entirely consistent with a negative reduced-form effect of a mass layoff itself (e.g., [Jacobson, LaLonde, and Sullivan 1993](#)).

profession has settled on a value for this elasticity close to zero.”⁷ Consistent with this view, we estimate a small earnings elasticity when taking a conventional quasi-experimental approach.⁸ We argue that such micro elasticities are uninformative of long-run responses among top earners, most of whom work in salaried jobs. Such workers cannot freely adjust earnings within a given job cell. They may change effort, but the earnings implications of changed effort play out dynamically and are often tied to job switches. Using hours worked as the outcome variable is not a solution because, for salaried workers, hours is a very limited measure of true effort.⁹ Rather, our proposed solution is to restrict attention to job switchers, maintaining earnings as the outcome variable. While some papers have studied heterogeneity in tax elasticities by job switching status (e.g., [Tortarolo, Cruces, and Castillo 2020](#)), we are not aware of any work that develops a theoretical framework and empirical approach using switchers to estimate welfare-relevant, long-run earnings elasticities.¹⁰

Our paper presents a new attempt to reconcile micro and macro evidence on labor supply. The micro-macro debate has focused on three issues: extensive margin responses ([Chetty, Guren, Manoli, and Weber 2013](#)), optimization frictions ([Chetty 2012](#)), and human capital accumulation ([Imai and Keane 2004](#); [Keane 2011](#); [Keane and Rogerson 2015](#)). Our approach is related to models incorporating human capital effects of effort — a specific channel through which dynamic returns may arise — but is at the same time fundamentally different. In standard human capital models, worker compensation is aligned with actual effort and productivity at any point in time, where productivity is allowed to change over time due to learning-by-doing or on-the-job training. Such effects are presumably slow-moving, and there is no role for discrete changes in earnings around job switches. Our approach using short-run switchers is not plausibly driven by human capital effects, while the human capital literature does not capture the effects studied here. Although

⁷To be clear, this assessment pertains to elasticities of *real* labor supply (or wage earnings) along the *intensive margin*, consistent with the focus of our study. Estimates of taxable income elasticities — including avoidance and evasion responses — can be considerably larger depending on the tax code and enforcement system. Estimates of extensive margin elasticities feature much greater variation across studies and less of a consensus (see [Kleven 2023](#)).

⁸Studying the same Danish tax reform, [Kreiner, Leth-Petersen, and Skov \(2016\)](#), [Jakobsen and Sogaard \(2022\)](#), [Lambanica and Pozzoli \(2022\)](#), and [Sigaard \(2022\)](#) also estimate small micro elasticities.

⁹In fact, this is one of the main reasons why the modern public finance literature has shifted its focus from hours-of-work elasticities to earnings elasticities. But by doing so, researchers solved one problem (the fact that hours responses are too narrow) by introducing another one (the fact that earnings responses are dynamic and delayed).

¹⁰A key advantage of our approach is that it relies on a widely available source of tax variation. This allows for applying the approach across different countries, reforms, and samples. A different approach would be to hunt for the “perfect” experiment to estimate long-run earnings responses. A compelling example of such an approach is provided by [Saez, Matsaganis, and Tsakloglou \(2012\)](#). They consider a cohort-based payroll tax reform in Greece and find zero long-run earnings responses. We are motivated by the view that such experiments are rare and context-specific. Building a consensus regarding the magnitude of long-run elasticities requires an approach that can be implemented repeatedly across settings and over time.

we argue that the long-run elasticity is larger than typical micro estimates, our estimates remain considerably smaller than those implied by a number of macro studies.¹¹

Our work is also related to the literature studying how optimization frictions shape observed labor supply. This includes a labor literature on hours constraints and adjustment costs (Altonji and Paxson 1986, 1988; Lachowska, Mas, Saggio, and Woodbury 2023) and a public finance literature showing that micro elasticities may be strongly attenuated by frictions (Chetty, Friedman, Olsen, and Pistaferri 2011; Chetty, Friedman, and Saez 2013; Kleven and Waseem 2013; Kleven 2016; Kreiner, Munch, and Whitta-Jacobsen 2015; Labanca and Pozzoli 2022; Anagol, Davids, Lockwood, and Ramadorai 2022). While dynamic returns to effort represent a conceptually different mechanism, their existence may be driven by an underlying information friction: the fact that the verification of effort and productivity is costly to employers. As we show, such verification costs give rise to an equilibrium with intermittent performance evaluations and dynamic returns. This insight is related to career-concern models (Harris and Holmström 1982; Holmström 1999) in which employers have imperfect information about worker productivity, and implicit contracts link current effort to future wages. Our model captures similar ideas in a simple manner and informs empirical work on labor supply responses.

Finally, our paper is linked to a large body of empirical work studying wage determination and careers. This includes papers that compare the implications of standard labor supply models and contract models for changes in earnings and hours over time (Abowd and Card 1987, 1989), arguing that the standard model fits the data poorly. It also includes papers that document the importance of job-to-job mobility for wage growth (e.g., Topel and Ward 1992; Farber 1999; Abowd, Kramarz, and Margolis 1999; Pavan 2011; Card, Heining, and Kline 2013; Card, Cardoso, and Kline 2016; Card, Cardoso, Heining, and Kline 2018). Consistent with these literatures, we take a contract view on employment relationships and emphasize the critical role of job switches for earnings dynamics. Given our empirical approach uses firm switchers, it is natural to ask if our earnings elasticities are mediated by firm-specific wage premia as estimated in the literature on AKM models (Abowd, Kramarz, and Margolis 1999). In other words, while our quasi-experimental estimates should be interpreted as worker responses (as they are based on tax variation across workers, not

¹¹Our agenda is also complementary to a paper by Scheuer and Werning (2017) on the optimal taxation of superstars. They argue that the welfare-relevant earnings elasticity in a superstar market is larger than in a standard labor market due to a job switching mechanism. When superstar workers are induced to provide greater effort through lower taxes, they anticipate being reassigned to a better job and this amplifies the incentive. We share the focus on job switching, but our model is otherwise different and highlights the importance of job switching effects for all salaried workers, not just superstars.

firms), these responses may be driven by workers moving into higher-wage firms. We estimate an AKM model to investigate this point, showing that the earnings responses are not driven by such effects of taxes on worker sorting across firms. This is consistent with our theoretical framework.

The paper proceeds as follows. Section 2 develops our theoretical model of dynamic compensation. Section 3 describes the data. Section 4 presents descriptive evidence on earnings and hours-worked patterns over the lifecycle, verifying the predictions of the model. Section 5 presents quasi-experimental evidence on earnings elasticities, using job switchers to uncover the long-run macro elasticity. Section 6 concludes and discusses policy implications.

2 A Theoretical Model of Dynamic Compensation

2.1 Setting

We consider a population of infinitely-lived workers with heterogeneous and time-varying productivities n_t . In each period, workers derive utility from consumption (which depends on *realized* earnings z_t) and disutility from effort (which depends on *latent* earnings y_t), where realized and latent earnings may be misaligned due to dynamic returns to effort. Flow utility is specified as

$$u_t = (1 - \tau) z_t - n_t v(y_t/n_t), \tag{1}$$

where τ is the marginal tax rate. The productivity parameter is specified as $n_t = g(t) + \mu$, where $g(t)$ is a common, deterministic lifecycle component and μ is an idiosyncratic, random shock.

The quasi-linear utility specification in (1) is commonly used in the literature on income taxation (e.g., Diamond 1998; Kleven, Kreiner, and Saez 2009; Saez 2010). However, the literature has focused on the standard labor supply model where $z_t = y_t$, i.e. where effort choices translate immediately and frictionlessly into realized earnings. In this special case, assuming that $v(x)$ takes the isoelastic form $\frac{\eta}{\eta+1} x^{\frac{\eta+1}{\eta}}$, worker optimization gives the familiar expression $z_t = y_t = (1 - \tau)^\eta n_t$, where η is the elasticity of earnings with respect to the marginal net-of-tax rate and productivity n_t represents potential earnings at a tax rate of zero.

We relax the assumption that effort maps immediately into earnings. In our model, realized earnings z_t change only at job events (such as switches between occupations or firms), which occur with probability λ in any given period. These job events align realized earnings with latent earnings (effort). Hence, we have

$$z_t = \begin{cases} y_t & \text{with probability } \lambda \\ z_{t-1} & \text{with probability } 1 - \lambda. \end{cases} \quad (2)$$

The basic idea is that worker effort is unobservable to employers absent costly performance evaluations, resulting in an employment contract where effort is rewarded discretely and intermittently at job events (performance evaluations). We start by assuming that the switching probability λ is exogenous, but we later develop a generalization in which the switching probability is endogenously set by firms facing effort verification costs. In either case, the value of λ determines the degree to which the return to effort is dynamic. The special case of $\lambda = 1$ corresponds to the standard labor supply model in which the return to effort is immediate. Conversely, if λ is small, the return to effort materializes far into the future in expectation.¹²

In this model, effort y_t is a choice variable and earnings z_t is a state variable. At time t , workers know z_{t-1} and n_t , and maximize expected lifetime utility with respect to current and future efforts. Denoting the discount factor by δ , the optimization problem can be written as

$$\max_{\{y_s\}_t^\infty} \sum_{s=t}^{\infty} \delta^{s-t} \mathbb{E} [u_s | z_{t-1}, n_t], \quad (3)$$

subject to the earnings dynamics in (2). The solution can be characterized as follows:

Proposition 1 (Optimal Effort). *Assuming $v(x) = \frac{\eta}{\eta+1} x^{\frac{\eta+1}{\eta}}$, the optimal choice of latent earnings (effort) is given by*

$$y_t = \left(\frac{\lambda}{1 - (1 - \lambda)\delta} \cdot (1 - \tau) \right)^\eta n_t \quad \forall t, \quad (4)$$

where η is the Hicksian elasticity of effort with respect to the marginal net-of-tax rate $1 - \tau$.

Proof. See Appendix B.1. ■

Optimal effort takes the standard form except for the adjustment factor $\frac{\lambda}{1 - (1 - \lambda)\delta}$, which captures the effect of dynamic returns to effort. When $\lambda = 1$, the level of effort is the same as in the standard model. Introducing dynamic returns ($\lambda < 1$) has two counteracting effects on the level of

¹²It is worth pointing out that our model is conceptually related to a large macro literature studying rigid prices and wages. This literature has developed models with time-dependent price adjustment rules (Taylor 1980; Calvo 1983), state-dependent price adjustment rules (Caplin and Spulber 1987; Caplin and Leahy 1991; Caballero and Engel 1991), and a combination of the two elements (Nakamura and Steinsson 2010). The earnings specification in (2) is a form of Calvo contract in which there is a constant probability that the earnings of a given worker are adjusted, independently of the time since the last adjustment. Importantly, our objective — understanding how earnings and effort respond to taxes — is fundamentally different from the focus in macroeconomics on nominal price rigidity and the impact of monetary policy.

effort. On the one hand, a lower λ implies that effort returns have longer delays and this disincentivizes effort (reflected in the numerator of the adjustment factor). On the other hand, a lower λ implies that effort returns, when they do materialize, are expected to last a longer and this incentivizes effort (reflected in the denominator of the adjustment factor). In the special case of $\delta = 1$, these two effects offset exactly and the level of effort is the same as in the standard model. Importantly, this point pertains to the level of *latent* earnings, whereas we are ultimately interested in the response of *realized* earnings to taxes. Even when $\delta = 1$, the model has very different predictions than the standard model. In fact, none of the key results derived below depend on the size of the discount factor.

Using equation (2), we can write average earnings at time t as a function of the average levels of effort from time 0 and the initial level of average earnings at time 0. We have

$$\bar{z}_t = \lambda \sum_{s=0}^t (1 - \lambda)^s \bar{y}_{t-s} + \left(1 - \lambda \sum_{s=0}^t (1 - \lambda)^s \right) \bar{z}_{-1}, \quad (5)$$

where average earnings \bar{z}_t equals a weighted average of historical efforts $\bar{y}_0, \dots, \bar{y}_t$ and initial average earnings \bar{z}_{-1} , with weights that depend on λ . This equation also describes each individual's expectation at time 0 of earnings at time t .

To summarize, the model has the following predictions on the variation in effort and realized earnings over time:

Proposition 2 (Effort and Earnings Predictions). *The dynamic compensation model ($\lambda < 1$) has the following predictions that differ from the standard labor supply model ($\lambda = 1$):*

1. *Average earnings \bar{z}_t depend on past effort choices $\bar{y}_{s < t}$, conditional on current effort \bar{y}_t .*
2. *The contemporaneous correlation between earnings z_t and effort y_t equals the per-period switching probability λ .*
3. *For each worker, the lifecycle profile of earnings z_t is discrete around job switches.*
4. *For each worker, the lifecycle profile of effort y_t is smooth around job switches, as long as productivity and taxes are smooth.*

Proof. (1) This follows from equation (5) derived above. (2) See Appendix B.2. (3) This follows directly from the specification in (2). (4) This follows from equation (4) in Proposition 1. ■

2.2 Earnings Responses to Taxes

The dynamic compensation model has important implications for earnings responses to tax policy and welfare measurement. To see this, consider a permanent change in the tax rate from time 0, assuming that the economy is initially in a steady state with constant average earnings, $\bar{z}_t = \bar{z}$.¹³ The welfare effect of such a tax change can be understood by considering its effects on tax revenue, $R_t = \tau \bar{z}_t$. A change in the tax rate has a mechanical effect on revenue, $dM_t = d\tau \cdot \bar{z}_t$, and a behavioral effect on revenue, $dB_t = \tau \cdot d\bar{z}_t$. Defining the elasticity of earnings at time t with respect to the net-of-tax rate as $\varepsilon_t^z \equiv \frac{d\bar{z}_t/\bar{z}_t}{d(1-\tau)/(1-\tau)}$, the ratio of behavioral and mechanical effects can be written as

$$dD_t \equiv dB_t/dM_t = \frac{\tau}{1-\tau} \cdot \varepsilon_t^z. \quad (6)$$

This is a standard formula for the marginal deadweight loss of taxation (see e.g., [Saez, Slemrod, and Giertz 2012](#); [Kleven 2021](#)), where the earnings elasticity ε_t^z is a sufficient statistic, conditional on τ . The issue is that, with dynamic compensation, ε_t^z increases over time and the measured welfare effect dD_t therefore depends on the time horizon of the elasticity estimation. Most quasi-experimental approaches allow only for the estimation of short-run elasticities and welfare effects, but policy design depends on long-run (steady state) welfare effects. In fact, assuming that the social planner puts equal weights on welfare now and in the future, the present value of social welfare is equivalent to steady state welfare.¹⁴

Computing the long-run welfare effect, dD_∞ , requires information about the long-run earnings elasticity, ε_∞^z . To show how such information might be obtained empirically, we derive the following properties of earnings elasticities.

Proposition 3 (Earnings Elasticities). *Consider a permanent change in τ from time $t = 0$, assuming that the economy is initially in a steady state with constant average earnings, $\bar{z}_t = \bar{z}$. In this case, we have*

1. *The long-run elasticity of realized earnings with respect to the net-of-tax rate equals $\varepsilon_\infty^z = \eta$, i.e. the elasticity of latent earnings (effort) characterized in equation (4).*
2. *The elasticity of realized earnings at time t is a downward-biased estimate of the long-run elasticity.*

¹³This assumption implies that we disregard any systematic lifecycle trends in earnings (i.e., $g(t)$ is constant). This simplifies the analysis, but is not important for the substance of the results. We consider the general case in an appendix, as discussed below.

¹⁴See Appendix B.3 for a proof.

Specifically,

$$\varepsilon_t^z = \alpha_t \eta \text{ where } \alpha_t = \lambda \sum_{s=0}^t (1-\lambda)^s \leq 1, \quad (7)$$

such that the elasticity starts at the short-run level $\varepsilon_0^z = \lambda \eta$ and increases gradually towards its long-run level $\varepsilon_\infty^z = \eta$.

3. For workers experiencing their first post-reform job switch at time t , the elasticity of realized earnings reveals the long-run elasticity, i.e. $\varepsilon_t^z|_{J_t=1} = \eta$ where $J_t = 1$ is an indicator for having the first post-reform job switch at time $t \geq 0$.

Proof. (1) This follows from equation (7) as $\alpha_\infty = 1$ regardless of λ . (2) Using that the initial steady state must have $\bar{z}_t = \bar{y}_t$, equation (5) implies that $\varepsilon_t^z = \lambda \sum_{s=0}^t (1-\lambda)^s \cdot \varepsilon_t^y$ where $\varepsilon_t^y \equiv \frac{d\bar{y}_t/\bar{y}_t}{d(1-\tau)/(1-\tau)}$. We have $\varepsilon_t^y = \eta$ from equation (4), which gives the relationship in (7). (3) It follows directly from equation (2) that $z_t|_{J_t=1} = y_t$ and, therefore, $\varepsilon_t^z|_{J_t=1} = \eta$. ■

This proposition shows that, in a world with dynamic compensation ($\lambda < 1$), estimates of short-run elasticities ε_t^z underestimate the welfare-relevant, long-run elasticity η . The bias is increasing in the degree of dynamic compensation (inversely related to λ). However, the last part of the proposition shows that the long-run elasticity can be uncovered from short-run responses by restricting the sample to job switchers because, for such individuals, realized and latent earnings momentarily coincide. The next section investigates identification based on switchers in greater depth.

Finally, while the derivations above disregarded lifecycle trends in earnings ($g(t)$ was assumed to be constant), the results are generalized to allow for such lifecycle dynamics in Appendix B.4. There we show that the formula for α_t becomes more involved, but it remains the case that it increases over time from $\alpha_0 = \lambda$ to $\alpha_\infty = 1$.

2.3 Heterogeneity and Identification

The preceding analysis allowed for heterogeneity in earnings via the idiosyncratic productivity term μ , but all other parameters of the model were assumed to be homogeneous across workers. Realistically, there will also be heterogeneity in effort elasticities η and switching probabilities λ . Denoting the joint density of these two parameters by $f(\eta, \lambda)$, we are interested in estimating the average long-run earnings elasticity, i.e.

$$\mathbb{E}[\varepsilon_\infty^z] = \int_\lambda \int_\eta \eta f(\eta, \lambda) d\eta d\lambda = \mathbb{E}[\eta]. \quad (8)$$

Our ability to identify this macro elasticity using switchers will depend on the properties of $f(\eta, \lambda)$. We can estimate the average earnings elasticity among workers making their first post-reform job switch at time t , i.e.

$$\mathbb{E}[\varepsilon_t^z | J_t = 1] = \int_{\lambda} \int_{\eta} \eta f(\eta, \lambda | J_t = 1) d\eta d\lambda, \quad (9)$$

where $f(\eta, \lambda | J_t = 1)$ denotes the density of η, λ among such first-time switchers. The following proposition characterizes the conditions under which this estimand recovers the long-run parameter of interest.

Proposition 4 (Identification). *Consider a permanent change in τ from time $t = 0$. For workers making their first post-reform job switch at time $t \geq 0$, the average elasticity of realized earnings identifies*

$$\mathbb{E}[\varepsilon_t^z | J_t = 1] = \mathbb{E}[\eta] + \frac{\text{cov}(\eta, \lambda(1-\lambda)^t)}{\mathbb{E}[\lambda(1-\lambda)^t]}, \quad (10)$$

where $\lambda(1-\lambda)^t \equiv P(J_t = 1 | \lambda)$ is the probability of making the first post-reform job switch at time t for a worker of type λ . If $\eta \perp \lambda$, we have $\text{cov}(\eta, \lambda(1-\lambda)^t) = 0$ and therefore $\mathbb{E}[\varepsilon_t^z | J_t = 1] = \mathbb{E}[\eta]$ for $\forall t \geq 0$.

Proof. From Bayes' Rule, we have $f(\eta, \lambda | J_t = 1) = \frac{P(J_t=1|\lambda)f(\eta,\lambda)}{P(J_t=1)}$, where $P(J_t = 1 | \lambda) = \lambda(1-\lambda)^t$ and $P(J_t = 1) = \mathbb{E}[\lambda(1-\lambda)^t]$ denote conditional and unconditional probabilities of making the first post-reform job switch at time $t \geq 0$. Inserting this into equation (9), we obtain

$$\mathbb{E}[\varepsilon_t^z | J_t = 1] = \frac{\mathbb{E}[\eta \cdot \lambda(1-\lambda)^t]}{\mathbb{E}[\lambda(1-\lambda)^t]}. \quad (11)$$

Using the definition of covariance ($\text{cov}(X, Y) = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$), this corresponds to the result in equation (10). ■

Under orthogonality between effort elasticities η and switching probabilities λ , the macro elasticity can be point identified. Importantly, this case is associated with an observable signature in the data, namely that the switcher elasticity $\mathbb{E}[\varepsilon_t^z | J_t = 1]$ is constant in t . We verify that this condition is satisfied in our empirical application, consistent with point identification. At the same time, because the condition may not hold across all settings, it is relevant to consider situations where η and λ are correlated. In such situations, partial identification is still possible. We have:

Corollary 1 (Partial Identification). *The probability of making the first post-reform job switch at time t , $P(J_t = 1 | \lambda) = \lambda(1-\lambda)^t$, is increasing in λ for $t < \frac{1-\lambda}{\lambda}$ and decreasing in λ for $t > \frac{1-\lambda}{\lambda}$. Therefore, if*

$\text{cov}(\eta, \lambda) > 0$, we have that $\text{cov}(\eta, \lambda(1-\lambda)^t)$ is positive at $t = 0$ and turns negative at a sufficiently large t . From equation (10), this implies that short-run switcher elasticities, $\mathbb{E}[\varepsilon_t^z | J_t = 1]$ for small t , provide upward-biased estimates of the average long-run elasticity $\mathbb{E}[\eta]$. In this case, with estimates of short-run switcher elasticities for periods $t = 0, \dots, T$, a lower bound is given by

$$\sum_{t=0}^T \Lambda_t \mathbb{E}[\varepsilon_t^z | J_t = 1] = \sum_{t=0}^T \Lambda_t \mathbb{E}[\eta | J_t = 1] \leq \sum_{t=0}^{\infty} \Lambda_t \mathbb{E}[\eta | J_t = 1] = \mathbb{E}[\eta], \quad (12)$$

where Λ_t denotes the share of workers making their first post-reform job switch at time t . Conversely, if $\text{cov}(\eta, \lambda) < 0$, then $\text{cov}(\eta, \lambda(1-\lambda)^t)$ is negative at $t = 0$ and turns positive at a sufficiently large t . From equation (10), this implies that $\mathbb{E}[\varepsilon_0^z | J_0 = 1] < \mathbb{E}[\eta]$ is a lower bound.

Proof. These results follow from Proposition 4 by noting that $\frac{dP(J_t=1|\lambda)}{d\lambda} = (1-\lambda)^t \left(1 - t \frac{\lambda}{1-\lambda}\right)$. ■

To conclude, in a world with dynamic compensation, short-run switchers can be used to either point identify or partially identify the long-run macro elasticity of interest.

2.4 Competitive Labor Market Equilibrium

With the simple wage dynamics in equation (2), if worker productivity n_t grows over time, output y_t is on average higher than wages z_t . This is inconsistent with a competitive equilibrium in which firms make zero expected profits. In this section, we show that the main results generalize to a competitive equilibrium where average wages equal average output. For simplicity, we assume that productivity grows at a constant rate g . If a worker experiences a job switch at time t , the wage of the worker is adjusted to the competitive market level \hat{z}_t and stays fixed until the next job switch. This implies that equation (2) is replaced by

$$z_t = \begin{cases} \hat{z}_t & \text{with probability } \lambda \\ z_{t-1} & \text{with probability } 1 - \lambda. \end{cases} \quad (13)$$

In this framework, we obtain the following result.

Proposition 5 (Competitive Equilibrium). *With perfect labor market competition and constant productivity growth g , the equilibrium level of latent earnings (effort) equal*

$$y_t = \left(\frac{\lambda}{1 - \delta(1-\lambda)(1+g)} \cdot (1-\tau) \right)^\eta n_t, \quad (14)$$

while the realized earnings of workers switching jobs at time t equal

$$\hat{z}_t = y_t \frac{1 - \delta(1 - \lambda)}{1 - \delta(1 + g)(1 - \lambda)}. \quad (15)$$

Proof. See Appendix B.5. ■

This proposition shows that, without productivity growth ($g = 0$), allowing for labor market competition and zero expected profits has no impact on effort and earnings. With positive growth, the earnings level of job switchers is higher than their contemporaneous effort (i.e., $\hat{z}_t > y_t$) to reflect that their output is expected to increase over time. This also leads to a higher level of effort y_t . Importantly, these level effects on effort and earnings do not change any of the main results. In particular, we see from Proposition 5 that the structural elasticity of effort is still equal to η , and that a given percentage change in effort still translates into the same percentage change in realized earnings among job switchers. These are the key properties needed for the remaining results to go through.

2.5 Endogenous λ

Appendix B.6 develops an extension of our model with an endogenous switching probability λ . In this model, worker effort is unobservable without a performance evaluation. The cost of evaluating a given worker is q and reveals true effort in the current period. Evaluations are carried out randomly with frequency λ . The equilibrium is given by the constrained-efficient solution in which chosen effort and evaluation frequency maximize worker-firm surplus. In a steady state, this amounts to maximizing

$$S = (1 - \tau) [y - q\lambda] - nv(y/n), \quad (16)$$

where we assume that evaluation costs $q\lambda$ are tax deductible. This will be the case if, for example, the costs of performance evaluations reflect verifiable labor or equipment costs.

In this framework, we obtain the following proposition.

Proposition 6 (Endogenous λ). *With a positive and finite evaluation cost q , the equilibrium evaluation frequency $\lambda \in (0, 1)$. The limit case of perfect verification ($\lambda = 1$ as in the standard model) is obtained for $q = 0$, while the limit case of no verification ($\lambda = 0$) is obtained for $q = \infty$. Outside these limit cases,*

we have that λ is decreasing in the evaluation cost q , increasing in the effort elasticity η , decreasing in the discount factor δ , and independent of τ .

Proof. See Appendix B.6. ■

This proposition microfound the dynamic compensation model and implies that all of our previous results generalize. Two specific points are worth highlighting. First, the switching probability is independent of the tax rate due to the assumption that evaluation costs are tax deductible. With partial or no deductibility, there would be an impact of taxes on the probability of switching. We will directly test for this in our quasi-experimental analysis. Second, in the extended model, heterogeneity in λ would be driven by underlying heterogeneity in evaluation costs q . The identification results of the previous section depend on the correlation between η and λ , which would be driven by the joint density of (η, q) and the fact that λ is directly increasing η by Proposition 6. Due to the direct effect, η and λ would tend to be positively correlated, in which case short-run switcher elasticities partially identify the long-run elasticity, with the lower bound given in equation (12).

3 Data

The empirical analysis is based on administrative data covering the full population of Denmark from 1980 to 2018. The data combine several administrative registers (linked at the individual level via personal identification numbers) and contain detailed information on earnings, hours worked, occupation, firm, and demographic variables. Virtually all of the information in the data is third-party reported (see [Kleven, Knudsen, Kreiner, Pedersen, and Saez 2011](#)).

Two features of the data are worth highlighting. First, the data are employer-employee matched and include detailed occupation codes, allowing us to observe jobs (firm \times occupation cells) at a granular level.¹⁵ The occupation codes build on the International Standard Classification of Occupations (ISCO), adapted by Statistics Denmark and called DISCO codes. The classification system has changed over time. Since 2010, occupations have been coded according to the DISCO-08 classification (563 occupations), while between 1991-2009, occupations were coded according to the DISCO-88 classification (372 occupations). We bridge this data break using a crosswalk developed by [Humlum \(2021\)](#). Prior to 1991, occupation codes were based on an older Danish classification

¹⁵Firm and occupation codes may sometimes change without an actual job switch (e.g., due to changes in ownership structure or reclassifications of worker groups). To avoid attributing such data changes to job switches, we drop all switches where more than 50% of an individual's coworkers move to the same job cell in the same year.

system (299 occupations). As this classification is still available after 1991, we are able to bridge the old occupation codes with the more recent ones. As a rule, private employers with at least 10 workers and all public employers must register and report the occupation of each worker to Statistics Denmark. For the remaining workers, Statistics Denmark imputes occupation codes based on industry, labor union, and education. Table A.1 in the appendix shows examples of top and bottom occupation titles, ranked by average wage earnings.

Second, the data include two administrative measures of hours worked. We are ultimately interested in the role of dynamic returns to *effort*, which is conceptually different from time spent at work.¹⁶ Still, given hours worked is a component of effort, we provide descriptive evidence on the relationship between hours and earnings that speaks to the predictions of the model. The first measure comes from a mandated pension scheme — *Arbejdsmarkedets Tillægspension* (ATP) — which requires employers to contribute on behalf of their employees based on individual working hours. The pension contribution is a function of a binned measure of working hours. Specifically, for someone paid monthly — the typical contract for salaried workers — the annual contribution depends on annual hours $\sum_{m=1}^{12} h_m$, where monthly hours h_m is divided into four bins.¹⁷ This measure is available for the entire period 1980-2018, but has the disadvantage of being capped at full time for all 12 months of the year.¹⁸ The second measure is better, but is only available since 2008. This measure provides information on uncapped hours for all workers at the monthly level. We use the first measure for analyses requiring a long panel and the second measure for analyses requiring us to capture hours variation precisely, including among full-time workers.

4 Descriptive Evidence on Dynamic Compensation

4.1 Four Descriptive Facts

To validate our theoretical model of dynamic compensation, this section presents descriptive facts that speak to the predictions of the model. The analysis leverages the granularity and statistical

¹⁶In general, observed working hours may deviate from true effort for two reasons. One reason is that working hours reported in the data may reflect contracted hours rather than actual hours, or some mixture between the two. The other reason is that unobserved effort choices influence the quality-adjusted hours relevant for earnings progression.

¹⁷This gives a total of 37 hours bins ($= 4 \times 12 - 12 + 1$) over a year.

¹⁸Even so, Figure A.1 in the appendix shows that the administrative pension measure of hours worked aligns well with information from survey data. The figure compares the relationship between earnings and hours worked in the administrative data (Panel A) to the relationship in labor force survey data (Panel B). The survey data contain information on self-reported hours worked (actual, uncapped hours). Reassuringly, the earnings-hours relationship is similar in the two datasets. But the survey measure is much more noisy than the administrative measure, especially at the top of the hours and earnings distribution, which is a key reason for using administrative data.

power of the Danish data to provide particularly clear evidence. As we shall see, the evidence is consistent with the dynamic compensation model *and* inconsistent with the standard labor supply model. The findings lend support to the quasi-experimental approach developed later.

4.1.1 Fact 1: Past Hours Worked Predict Current Earnings, Conditional on Current Hours

Any model of dynamic returns to effort predicts that earnings depend on past hours worked, even after conditioning on current hours worked. Figure 1 investigates if this prediction is supported by the data. The figure is based on a balanced panel of workers observed between the ages of 20 and 50, showing how earnings at age 50 relate to current and past hours worked.¹⁹ Each panel shows the non-parametric relationship between the average earnings rank at age 50 and hours worked at different ages. Panel A considers current working hours (at age 50), while Panel B considers past working hours (at ages 40-49). There is a strong positive relationship in both panels: working more hours, past or present, is associated with higher earnings. However, the fact that hours worked are correlated over time complicates the interpretation. The positive relationship between earnings and past hours worked may reflect that variation in past hours captures variation in current hours. In addition, variation in past hours may be correlated with variation in productivity parameters that impact earnings directly. The subsequent panels investigate if the predictive power of past hours is the result of such correlations.

Panel C shows the relationship between earnings rank at age 50 and hours worked between ages 40-49 without any controls (blue dots), with controls for current hours (orange dots), and with controls for both current hours and demographic variables that proxy for earnings capacity (red dots).²⁰ The controls dampen the correlation between earnings and past hours as one would expect, but the relationship remains strong. The expected earnings rank at age 50, conditional on current hours and demographics, increases from the 25th percentile to the 65th percentile as annual hours worked over the preceding 10 years increases from zero to 2,000 hours. The relationship becomes stronger at high levels of hours and earnings: making it to the top of the distribution requires consistently high effort over time. As a robustness check, Panel D considers the effect of past hours worked over a longer time horizon. This hardly affects the relationship.

These findings suggest that the return to effort (hours) has a strong dynamic component. The

¹⁹In this analysis, because we consider a long panel, working hours are measured using employer pension contributions (ATP) available across all years of the data, as described in Section 3.

²⁰The demographic controls include dummies for education level (8 categories), gender (binary), children (binary), and marital status (7 categories).

evidence contradicts standard labor supply models in which current hours of work is a sufficient statistic for earnings.

4.1.2 Fact 2: Hours and Earnings Changes are Contemporaneously Unrelated at the Top

Another prediction of the dynamic compensation model is that the contemporaneous correlation between hours and earnings changes is small.²¹ This stands in contrast to standard models of hourly-paid workers in which the two are perfectly correlated. To shed light on the two models, Figure 2 plots changes in log hours against changes in log earnings for workers in different parts of the earnings distribution: the bottom 20%, the top 20%, the top 10%, and the top 1%. The average relationship in each segment is shown by blue dots, while examples of specific occupations are shown by red triangles and diamonds. To capture hours variation even among top earners, these graphs use monthly information on uncapped hours available since 2008.²² Importantly, the theoretical prediction is about the intensive margin, not the extensive margin. To avoid effects from the extensive margin, we consider average monthly outcomes over the year calculated across months with positive hours and earnings. Our results therefore capture correlations at the annual level, neutralizing any variation from the number of months worked.

As can be seen from the figure, there is a stark difference between workers at the bottom and the top. At the bottom, the relationship between log hours and log earnings is very close to the 45-degree line, as predicted by standard models of hourly-paid workers. This is natural when considering typical occupations at the bottom: the examples provided in the figure — cleaners and waiters — are hourly-paid jobs. At the top of the distribution, on the other hand, the relationship between log hours and log earnings is much flatter. The slope remains positive (albeit small) in the top 20% and top 10%, while it is virtually zero in the top 1%. Again, this is natural when considering typical occupations at the top. All of the examples shown — pharmacists, engineers, lawyers, CEOs, etc. — are salaried jobs that have no immediate link between working hours and earnings. The contemporaneous correlation is not quite zero, consistent with our theoretical model in which the probability of switching jobs creates a within-period link between effort and earnings. The fact that the correlation decreases as we move further into the top tail may be driven by heterogeneity in job switching probabilities across different quantiles of earnings. As workers reach the very top, switching across occupations or firms becomes less frequent.

²¹Recall that, in our model, the contemporaneous correlation coefficient between hours and earnings is equal to the unconditional job switching probability $\lambda \ll 1$.

²²See section 3 for details.

These findings suggest that, while the standard labor supply model may be a reasonable approximation for the bottom of the labor market, it is a poor description of the top of the labor market. For the top, the patterns are instead consistent with the dynamic compensation model.

4.1.3 Fact 3: Lifecycle Profiles in Earnings are Driven by Discrete Job Switches

We now turn to the role of job switching for the dynamics of earnings. The key idea of our paper is that the returns to effort are realized dynamically, at the point of job switches. Importantly, the objective in this section is not to investigate if job switches have causal effects on earnings. In our model, the causal driver of realized earnings is latent earnings, the profile of which reflects changes in effort and productivity over time. Job switches mediate the link between effort and earnings, but have no independent causal effect. Our objective is to verify if the prediction regarding the mediating role of job switches for earnings dynamics is borne out by the data. This is critical for understanding how to model labor supply and for developing an empirical strategy to estimate long-run responses to taxes.

Leveraging the granularity of the Danish data, we measure job switches as transitions between firm \times occupation cells. The first set of results is presented in Figure 3. Based on a balanced panel of workers observed between the ages of 20 and 50, the figure plots lifecycle profiles of earnings for different groups of workers. In Panel A, we compare workers in the top 10% and the bottom 50% of the earnings distribution at age 50. The two groups start at very similar earnings levels at age 20, but workers who make it to the top have a steeper lifecycle profile. The divergence in lifecycle profiles, illustrated in Panel B, leads to an earnings gap of about 1.7 log points at age 50. The question is how much of this divergence can be attributed to switches between job cells.

Panel C provides a striking answer. Starting from the raw difference in earnings profiles (dark blue), it shows the difference net of occupation fixed effects (light blue), net of occupation \times firm fixed effects (orange), and net of occupation \times firm \times individual fixed effects (red). The impact of each set of controls depends on the order in which we include them, but we are ultimately interested in the total impact of including all of them. Theoretically, the discreteness of the lifecycle profile at job switches is a within-worker phenomenon, which is why we interact job fixed effects with individual fixed effects. The evidence shows that, within individual, job fixed effects explain almost all of the divergence in the lifecycle profiles of top and bottom earners. Within job cells, there is virtually no divergence between the two groups. In Panel D, we move further into the top tail of the distribution, comparing top 1% earners to bottom 50% earners. The results are very

similar: 93% of the lifecycle divergence can be attributed to job transitions.

Do these results reflect that top earners make better switches or that they make more switches? Figure A.2 in the appendix shows that it is the former. The figure plots the distribution of the number of switches in the top and bottom samples analyzed above. The distributions are broadly similar in the different samples. The average number of switches is about 10 at the top and 9 at the bottom, corresponding to roughly one switch every three years. It is worth noting, however, that switching activity is not evenly spaced over the lifecycle. As workers age and reach higher earnings levels, switching becomes less frequent.

Another way of analyzing the importance of job switches is by decomposing the variance in earnings over the lifecycle into between-job variance and within-job variance. To implement this variance decomposition, note that the earnings of individual i in job j at time t can be written as

$$z_{ijt} = \bar{z}_{ij} + (z_{ijt} - \bar{z}_{ij}), \quad (17)$$

where \bar{z}_{ij} denotes the average earnings of individual i in job j . Demeaning by the average earnings of individual i , \bar{z}_i , and taking variances gives

$$\text{var}(z_{ijt} - \bar{z}_i) = \underbrace{\text{var}(\bar{z}_{ij} - \bar{z}_i)}_{\text{Between Job}} + \underbrace{\text{var}(z_{ijt} - \bar{z}_{ij})}_{\text{Within Job}}, \quad (18)$$

where we use that $\text{cov}(\bar{z}_{ij} - \bar{z}_i, z_{ijt} - \bar{z}_{ij}) = 0$. Hence, the variance of individual earnings over the lifecycle (relative to its mean) can be decomposed into between-job and within-job variances, with no covariance term.

Figure 4 presents the results of such an analysis. Considering the same panel of workers analyzed above, the figure plots the total variance of earnings (blue), the between-job variance (red), and the within-job variance (yellow) by earnings percentile at age 50. At all percentiles shown, virtually all of the lifecycle variation in earnings can be attributed to between-job variation. Consistent with the lifecycle graphs, between-job variance accounts for more than 90% of total variance. This holds at all earnings levels between the 50th and the 100th percentile. The robustness of the decomposition to the level of income suggests that this is a general feature of job contracts among salaried workers, who dominate a broad segment of the earnings distribution.

These results imply that job switches are central to understanding earnings dynamics and, by extension, to estimating earnings responses to taxes. Standard empirical approaches are not plau-

sible in a world where all of the action is concentrated at switches that happen only intermittently.

4.1.4 Fact 4: Earnings Increase Discretely at Promotions, with No Change in Hours

The last piece of descriptive evidence focuses on earnings and hours changes around promotion events. In our model, positive job events — events where latent earnings are higher than current earnings — lead to sharp increases in earnings, with no change in effort. We verify this prediction based on an event study of within-firm promotions. Defining promotion as a switch to an occupation cell in which median earnings are at least 10% higher, we compare the outcomes of promoted and unpromoted co-workers over time. The results are robust to considering other promotion thresholds such as 5% or 20%.

To conduct the analysis, we use monthly data on earnings and hours worked, aggregated to the quarterly level. We match each promoted worker to their unpromoted co-workers within the same firm, giving unpromoted workers the same event time.²³ Letting Y_{iq} be the outcome of individual i in quarter q , indexed such that $q = 0$ is the first quarter of promotion, the event study regression is specified as

$$Y_{iq} = \sum_j \alpha_j \cdot Event_{j=q} + \beta \cdot Treat_i + \sum_{j \neq -1} \gamma_j \cdot Event_{j=q} \cdot Treat_i + \phi_{q \in t} + \phi_a + \nu_{iq}, \quad (19)$$

where $Event_{j=q}$ is a quarterly event time dummy, $Treat_i$ is a promotion dummy, $\phi_{q \in t}$ is a calendar year fixed effect, and ϕ_a is an age fixed effect. We include year and age fixed effects to neutralize time and lifecycle trends unrelated to promotions. The coefficients of interest are γ_q . These are difference-in-differences coefficients that capture the effect of promotion in quarter q relative to the pre-promotion quarter $q = -1$ for promoted relative to unpromoted co-workers.

Figure 5 plots the difference-in-differences coefficients by event time for earnings (Panel A) and hours worked (Panel B). We see sharp effects on earnings: pre-trends are parallel, promoted workers experience a jump of about 4% at the time of the event, and the effect is stable over time. Conversely, there are no such effects on hours worked, which are smooth around the time of promotion. While these series have been normalized to zero at event time -1 , it should be noted that there are level differences between promoted and unpromoted workers. Workers who get promoted tend to have higher working hours and earnings leading up to the event, consistent with the idea that promotions reward past effort.

²³In selecting the sample, individuals are required to stay in the same firm from two quarters before promotion to two quarters after promotion.

Harking back to comments made in the previous section, these event studies should not be interpreted as estimating a causal effect of promotions on earnings. Viewed through the lens of our model, realized earnings are ultimately driven by effort and productivity, mediated through job switches due to the structure of job contracts. Saying that promotions are the reason for earnings jumps corresponds to saying, for example, that tenure decisions are the reason for changes in academic salaries. This is true only in a narrow sense. The real reason is the quality of the academic CV, the returns to which are materialized at discrete tenure events. Our model and evidence imply that this is a general feature of job contracts, where workers cannot influence earnings through effort without a discrete job event.

5 Estimating Earnings Elasticities with Dynamic Compensation

5.1 A Quasi-Experimental Approach Using Job Switchers

We now turn to the study of long-run earnings responses to taxes. To obtain exogenous variation in tax rates, we use a major tax reform implemented in Denmark in 2009-10.²⁴ Prior to the reform, income was taxed according to a progressive schedule with three brackets, commonly referred to as the bottom, middle, and top brackets. The 2009 reform eliminated the middle bracket and raised the top bracket threshold. The implication of these policy changes was that taxpayers above an income threshold experienced lower marginal tax rates, while those below the threshold were largely unaffected. The threshold that separates treatments and controls is located at around the 70th percentile of the income distribution. Figure 6 shows the evolution of marginal tax rates for taxpayers above and below the treatment threshold.²⁵ The reform-induced reduction in the top marginal tax rate was large, about 11 percentage points.²⁶ While the tax cut was relatively large, there is otherwise nothing unique about this experiment: it creates tax variation by income level of the sort typically used in the literature on behavioral responses to taxes (see [Saez, Slemrod, and Giertz 2012](#)). Indeed, the objective is to demonstrate our approach using a widely available source

²⁴We refer to [Jakobsen and Søgaard \(2022\)](#) for a detailed description of the Danish tax system and 2009 reform.

²⁵The figure depicts the *average* reduction in the marginal tax rate above the treatment threshold. The policy changes described above were associated with heterogeneity in tax rate reductions depending on location within the treated income interval.

²⁶To be exact, the average reduction in the marginal tax rate was 11.43pp above the threshold and 1.70pp below the threshold. The difference in tax treatment was therefore 9.73pp. These tax cuts had a large effect on the net-of-tax rate $1 - \tau$ relevant for elasticity calculations: the net-of-tax rate increased by 24% above the threshold and by 3% below the threshold, a treatment difference of 21pp.

of tax variation.²⁷

Standard quasi-experimental studies of earnings responses to tax reform compare treated and untreated workers in a short time window around a reform. As shown by our theoretical model, such approaches underestimate long-run responses in a setting where the returns to effort are dynamic and mediated by job switches. In general, the feasible time window in quasi-experimental studies is too short to allow most workers to switch jobs and realize the return to higher effort. The attenuation bias depends on the degree of job mobility; it will be stronger in populations where job switching is less frequent. Our theoretical model shows that the long-run earnings response can be point identified from short-run switchers if the timing of switching is not selected on the underlying structural elasticity, and that it can be partially identified in the presence of selection. The model also shows that the time profile of the observed earnings elasticity for switchers (within a short window) can be used to separate the cases with and without selection.

Our empirical strategy starts from a difference-in-differences specification of the form

$$\Delta \log z_{it} = \beta_0 + \sum_{b \neq b_0} \beta_b \cdot \mathbf{1}[y_{i0} \in b] + \sum_{b \neq b_0} \gamma_b \cdot \mathbf{1}[\text{switch}_{it}] \cdot \mathbf{1}[y_{i0} \in b] + \nu_{it}, \quad (20)$$

where z_{it} denotes the labor income of individual i at time t , y_{i0} denotes the taxable income of individual i prior to the reform, and $\mathbf{1}[\cdot]$ is an indicator function.²⁸ The measure of taxable income that determines treatment status includes labor income, transfers, pensions, alimony, and certain capital income items. We divide this taxable income measure into a set of discrete bins b , omitting a bin b_0 below the treatment threshold. We interact treatment status (pre-reform income bin) with an indicator for switching job, defined as moving between firm \times occupation cells.²⁹ This allows us to estimate difference-in-differences coefficients separately for stayers ($\hat{\beta}_b$ in income bin b) and movers ($\hat{\beta}_b + \hat{\gamma}_b$ in income bin b). By estimating these coefficients by income bin, we are able to investigate heterogeneity in behavioral responses across different income levels.

Given the specification assigns treatment status based on pre-reform income, the main threat to identification is the presence of non-tax effects on earnings growth $\Delta \log z_{it}$ that vary by pre-reform income level y_{i0} . The most obvious confounder is mean reversion, as discussed extensively in the

²⁷It should be noted that the tax reform happened around the time of the great recession, raising potential concerns about heterogeneous effects of the recession and recovery across treatments and controls. As we discuss later, our results are not consistent with bias from such macro confounders.

²⁸The outcome variable z_{it} equals total labor compensation, including contributions to employer-administrated retirement savings plans. This ensures that our estimates are not influenced by shifting between taxable compensation and tax-deferred retirement savings contributions.

²⁹Specifically, $\mathbf{1}[\text{switch}_{it}]$ equals one for workers who make their first post-reform job switch at time t .

literature (see [Saez, Slemrod, and Giertz 2012](#)). If income consists of both permanent and transitory components, those with high pre-reform incomes tend to be selected on positive transitory shocks, creating downward bias in the estimated responses to lower taxes as the transitory shocks dissipate over time. Following [Jakobsen and Sogaard \(2022\)](#), we address this issue by running our regression in pre-reform and post-reform datasets separately. The pre-reform specification considers earnings growth between 2006-08 by 2006 income bin. The resulting placebo estimates capture the effects of non-tax confounders assuming these are stable over time, which we verify in the data. The post-reform specification considers earnings growth between 2008-10 (and later) by 2008 income bin. Denoting the placebo estimates by superscript P , we estimate behavioral responses based on a triple-differences approach: $\hat{\beta}_b - \hat{\beta}_b^P$ for stayers and $\hat{\beta}_b + \hat{\gamma}_b - (\hat{\beta}_b^P + \hat{\gamma}_b^P)$ for movers. We also consider a quadruple-differences approach based on comparing the triple-differences estimates for movers and stayers, i.e. $\hat{\gamma}_b - \hat{\gamma}_b^P$.

We convert the estimates of earnings responses into elasticities with respect to $1 - \tau$. For this purpose, we pool the income bins in equation (20) into treated and untreated regions (above and below the threshold), re-estimating the equation with these broader treatment categories. When using the quadruple-differences specification, the elasticity can be calculated as

$$\varepsilon|_{\text{switch}} = \frac{\hat{\gamma} - \hat{\gamma}^P}{\mathbb{E} [\Delta \log (1 - \tau_{it}) | T] - \mathbb{E} [\Delta \log (1 - \tau_{it}) | C]}, \quad (21)$$

where $\hat{\gamma}$ and $\hat{\gamma}^P$ are pooled regression coefficients that capture the average effect on switchers above the treatment threshold. The denominator equals the average change in $\log (1 - \tau_{it})$ from before to after the reform for treatments relative to controls. Because the tax change in the control group is close to zero, heterogeneity in responsiveness is not an identification concern here.

In general, estimating behavioral responses by comparing individuals in different pre-reform tax brackets is associated with attenuation bias due to movements between brackets over time. Some taxpayers assigned to the treatment group will have incomes below the treatment threshold after the reform, and vice versa. As a result, the reduced-form estimates described above should be interpreted as intention-to-treat (ITT) effects. Conceptually, we are more interested in treatment-on-the-treated (TOT) effects. To estimate TOT elasticities, we specify the pooled version of equation (20) in terms of current taxable income y_{it} , constructing instruments based on pre-reform taxable income y_{i0} . The model is estimated using 2SLS. The elasticity is still defined as in equation (21), only the numerator is based on TOT coefficients obtained from the 2SLS estimation.

As discussed, for the switcher elasticity to point identify the long-run elasticity of interest, the timing of switching cannot be selected on the underlying structural elasticity. We provide several analyses to rule out such selection. First, we estimate if the switcher elasticity varies by the timing of switching jobs following the reform. We find that the elasticity is very stable over time, consistent with our theoretical prediction under orthogonality between switching probabilities and structural elasticities. The stability of the switcher elasticity also alleviates concerns about confounding effects of the great recession. Second, we investigate if switcher characteristics respond to the reform, using the same empirical design as we use for estimating earnings responses. That is, we ask if any observable differences between switchers above and below the treatment threshold change from before to after the reform. Looking at a wide range of characteristics, we do not find effects on any of them. As a result, controlling for demographic variables in specification (20) hardly affects the estimates. Third and finally, we restrict the sample to plausibly exogenous switches, namely those triggered by mass layoffs. We show that mass-layoff switchers feature similar earnings responses as the full sample of switchers.

5.2 Impact of Tax Reform: Switchers vs Non-Switchers

5.2.1 Reduced-Form Effects

We start by investigating the short-term earnings responses to the 2009 tax reform. Figure 7 shows results for all workers (Panel A) and for job movers vs job stayers (Panel B). Job movers are defined as workers who switch between firm \times occupation cells. The figure plots changes in log earnings between 2008-10 (actual experiment) and between 2006-08 (placebo experiment) by pre-reform income bin, depicting the income threshold that separates treatments and controls by a vertical line. As discussed in the previous section, the placebo series capture non-tax effects on earnings growth that vary by income level. This includes mean reversion: those with high pre-reform incomes tend to be selected on positive transitory shocks, reducing their earnings growth over time regardless of the tax cuts. Comparing the 2008-10 and 2006-08 series (above vs below the treatment threshold) gives a triple-differences estimate that controls for mean reversion.

Consider first the results for all workers in Panel A. The 2008-10 series shows that earnings growth is declining in pre-reform income level. As a result, a standard difference-in-differences approach that compares earnings growth above and below the treatment threshold yields a negative effect. As can be seen from the 2006-08 series, this reflects the aforementioned bias from

mean reversion. This series shows that, before the reform, there was a similar negative relationship between earnings growth and baseline income level. The two series track each other almost perfectly below the treatment threshold — consistent with mean reversion being stable over time — but diverge above the threshold. This divergence represents the causal effect of the tax reform on earnings. The effect is visually clear and precisely estimated, but modest in size. Reducing the top marginal tax rate by about 10pp increases earnings by less than 2%.

Consider now the results for job movers vs job stayers in Panel B. The modest response in the full sample masks strong heterogeneity by job switching status. For job movers, the pre-reform and post-reform series track each other perfectly below the treatment threshold and diverge sharply above the threshold. The earnings responses are large — about three times larger than in the full sample — and increasing in baseline income. For job stayers, on the other hand, we see no positive earnings responses to lower taxes. In fact, the post-reform series lies marginally below the pre-reform series above the treatment threshold, corresponding to a small negative effect. This seems to reflect an imperfection in the mean reversion adjustment for this subsample: the pre-reform series is rotated relative to the post-reform series, creating a small negative divergence in the treated range. The fact that the rotation occurs across both the treated and untreated regions suggests that the true response among job stayers is zero rather than negative.

The preceding analysis pool job movers regardless of whether they switch between firms or occupations. To see if the type of switch matters, Figure A.3 in the online appendix reproduces the analysis for firm switchers and occupation switchers separately.³⁰ The empirical patterns are very similar in these subsamples. These results suggest that the mediating role of job switches is independent of the particular type of switch. This is consistent with our theoretical model in which any switching event is associated with a performance evaluation that realigns latent and realized earnings.

The combination of large responses among job movers and no responses among job stayers suggests that the standard approach provides strongly downward-biased estimates of the true long-run effect. The reason is that job stayers do not stay in their current jobs permanently; they will eventually become job movers themselves and reveal their tax-induced earnings response at that time.³¹ The idea of our approach — using short-term movers to estimate long-term responses

³⁰Job switchers may change firm and occupation simultaneously. To retain statistical power, the subsamples in Figure A.3 include either all firm switches (even if occupation also changes) or all occupation switches (even if firm also changes). The results for within-occupation firm switches are similar to those shown in the figure, but the results for within-firm occupation switches are noisy as such switches represent a relatively small fraction of the data.

³¹As shown in Appendix Figure A.2, all workers switch jobs — and typically many times — over their careers.

in the population — is that job stayers increase *latent* earnings by as much as job movers, but that their return is not realized until the time of moving. As shown above, point identification of the long-run elasticity requires orthogonality between structural elasticities and switching probabilities. This is associated with an observable pattern in the data, namely that the earnings response by job movers is constant in the timing of switching.³² To investigate this point, we now consider the time profile of earnings responses among job movers.

To estimate a time profile of earnings responses that relates to our theoretical predictions, it is important to convert the intent-to-treat (ITT) effects into treatment-on-the-treated (TOT) effects. As discussed above, ITT estimates based on comparing individuals in different pre-reform tax brackets are associated with attenuation bias due to people moving across brackets after the reform. Such attenuation becomes stronger over time, thus confounding the interpretation of the time profile in ITT effects. We estimate TOT effects based on a 2SLS specification in which the actual post-reform tax bracket is instrumented using the pre-reform tax bracket. The details were provided in section 5.1.

Figure 8 shows the time profile of ITT effects (Panel A) and TOT effects (Panel B) on earnings. The red series provide estimates of long-run responses using our dynamic approach. Each dot represents a quadruple-differences estimate based on comparing earnings growth above vs below the treatment threshold for job movers vs job stayers in a post-period relative to a pre-period.³³ The pre-reform part of the series (2002-2008) gives placebo estimates by comparing, for each year t , earnings growth between year $t - 2$ and t to earnings growth between year $t - 4$ and $t - 2$. The post-reform part of the series (2009-2014) gives cumulative estimates by comparing, for each year t , earnings growth between 2008 and year t to earnings growth between 2006 and 2008. As can be seen from the figure, the pre-reform estimates are very close to zero — validating the empirical design — while the post-reform estimates are stable from year 2010 onwards. The estimates are smaller in 2009, which is natural given the reform was not yet fully phased in at that time. Leaving aside this partially treated year, the earnings responses among job movers do not depend on the timing of the switch. This pattern is consistent with point identification of the long-run earnings elasticity, as shown in section 2. Furthermore, the stability of earnings responses over time suggests against concerns about confounding effects of the great recession and recovery.

³²Conversely, if structural elasticities and switching probabilities are correlated, the earnings response by switchers is either declining (positive correlation) or increasing (negative correlation) as we consider switches farther removed from the time of the reform.

³³In this figure, we control for baseline job cell by including occupation×firm fixed effects. As we show later, the results are robust to adding or dropping controls.

Can we identify the true long-run elasticity using a standard approach simply by considering a sufficiently long post-reform period? To answer this question, Figure 8 provides time profiles of simulated earnings responses in the full sample of workers, including job stayers. These profiles are illustrated by the black series in the figure. They are based on scaling the estimates for job movers using the cumulated fraction of workers who have made a post-reform job switch, as characterized by equation (7). The earnings responses in the full sample increase over time as more workers make job transitions, but the convergence between the standard and dynamic approaches is slow. Even six years out, the standard estimate is only about 60% of the dynamic estimate. Most tax reform studies are unable to estimate six-year effects due to instability in bracket location (treatment status) over time. For example, the influential study by Gruber and Saez (2002), along with many studies that followed their approach, estimate effects over at most three years.

Finally, we ask if taxes impact the probability of switching between firms or occupations. As shown in section 2.5, even when switching probabilities are endogenously determined based on the costs of verifying true effort, they do not respond to taxes as long as verification costs are tax deductible. Figure 9 investigates if this prediction is borne out by the data. It shows switching probabilities for workers above and below the treatment threshold between 2002 and 2014. The treatment and control series track each other over time, consistent with no effect on the probability of switching jobs. The difference-in-differences estimate of the effect on switching equals -0.008 (0.001). While this estimate represents the effect on *any* switch between firms or occupations, Figure A.4 in the appendix investigates if there is heterogeneity by type of switch. It shows a precisely estimated zero effect for both firm and occupation switches. These findings, besides confirming one of the predictions of our model, further alleviates concerns that job switching is selected. If the tax reform had increased the likelihood of switching, we would expect the marginal switchers — presumably those with the strongest incentive to change jobs — to be selected on large elasticities. The absence of tax-induced job switching rules out such endogeneity.

5.2.2 Elasticities

In Table 1, we convert our estimates of earnings responses to the 2009 tax reform into elasticities with respect to the marginal net-of-tax rate. The table shows elasticities based on the dynamic approach (left panel) and the standard approach (right panel) over time. As described above, the dynamic approach relies on estimating earnings responses for job movers relative to job stayers, while the standard approach relies on estimating the average earnings response among all workers

(movers and stayers). The time series of earnings responses used in the elasticity calculations are those presented in Figure 8.³⁴ The table provides both intention-to-treat (ITT) and treatment-on-the-treated (TOT) estimates. The estimates of primary interest — highlighted by boldface in the table — are the TOT elasticities based on the dynamic approach.

The table provides three main insights. First, the dynamic approach using job movers yields sizable earnings elasticities, between 0.41 and 0.47. These are considerably larger than typical micro estimates of real labor supply elasticities at the intensive margin (Blundell and MaCurdy 1999; Saez, Slemrod, and Giertz 2012). Second, the elasticities for job movers are stable over time. As we have shown theoretically, this is critical for identification: the stability of elasticities among job switchers implies orthogonality between structural elasticities and switching probabilities, in which case the long-run elasticity is point identified. Third, the elasticities obtained from the standard approach are much smaller, but increasing over time as more workers make job transitions and realize their dynamic returns. The elasticity increases from 0.17 two years after the reform to 0.26 six years after the reform. The elasticity is still considerably lower than the dynamic long-run elasticity even after six years. The reason is that job switching is infrequent for a sizeable fraction of high-income workers, creating downward bias in the (standard) elasticity estimate even over long time horizons. As mentioned above, the six-year window considered here is longer than those considered in most quasi-experimental tax reform studies (Saez, Slemrod, and Giertz 2012).

Figure 10 investigates heterogeneity in elasticities by income level. It plots elasticities based on the standard approach (black series) and the dynamic approach (red series) above different income cutoffs, depicted on the x-axis. All elasticities are based on TOT effects between 2008-10. The figure shows that the standard elasticity is small and virtually constant in income, whereas the dynamic elasticity is large and increasing in income. The dynamic elasticity increases from 0.41 when considering all job movers above the treatment threshold (corresponding to the average elasticity in Table 1) to 0.60 when considering movers at the very top of the distribution.³⁵ The fact that the gap between dynamic and standard elasticities increases in income suggests that dynamic compensation effects become more important as we move further into the upper tail of the distribution.

³⁴In this figure, the standard approach is implemented by scaling the movers' estimate in each year t by the cumulated fraction of workers who have made a post-reform job move by year t . This implies that we assign a zero elasticity to job stayers consistent with the evidence discussed above.

³⁵The highest income cutoff shown in the figure (675,000 DKK) corresponds approximately to the 98th percentile of the income distribution.

5.3 Identification: Is Switching Selected?

The preceding analysis presented two findings that lend support to our identification strategy: earnings responses for job switchers are stable over time and switching probabilities do not respond to tax changes. As discussed, both of these results are consistent with the idea that short-term switcher elasticities can be used to uncover long-run elasticities. In this section, we present two additional identification checks, focusing on whether job switching is selected in a way that could bias our estimates.

The first check investigates if job switching is selected on observables given our empirical design. Specifically, treating demographic characteristics as dependent variables in our quasi-experimental analysis, we ask if any of them respond to the reform? The evidence is presented in Figure 11. It plots time series of different demographic characteristics for job movers above and below the treatment threshold. We consider six variables that are important for labor supply and earnings: age, gender, marital status, number of children, occupational rank, and firm size. Our measure of occupational rank is based on ordering occupation cells by their mean earnings. For all six variables, the treatment and control series are virtually parallel over time. There is no inflection point for treated job movers relative to untreated job movers after the reform. The difference-in-differences estimates of the effects on demographics (displayed in the figure) are very small, albeit statistically significant in several cases due to the statistical power of our data.

These findings suggest that adding control variables to our quasi-experimental specification should not make any material difference to the estimates. This is confirmed in Figure 12. The figure shows the time profile of estimated earnings responses for job movers under different sets of controls: no controls, demographic controls, job cell fixed effects, and finally demographics and job cell fixed effects combined. The time profile of treatment effects is very similar in the different specifications.

The second check investigates if job switching is selected on unobservables correlated with earnings growth. To address this question, we focus on switches triggered by a plausibly exogenous event: a mass layoff. A large literature on the effects of job displacement has used mass layoffs as a source of exogenous variation (e.g., [Jacobson, LaLonde, and Sullivan 1993](#)). Building on this literature, we implement our quasi-experimental approach in the sample of workers who switch firms following a mass layoff. The findings are presented in Figure 13. The figure plots earnings responses by income bin for the full sample of movers (Panel A) and for the sample of mass-layoff movers (Panel B). Mass-layoff movers are defined as workers who switch to a new

firm, coming from a firm that reduced their workforce by at least 30% in the year of the switch.³⁶ We find that the earnings responses to lower taxes are large even among mass-layoff movers. In fact, the responses are slightly larger than in the full sample of movers. Hence, the large switcher elasticities documented above do not appear to be driven by selection into switching.³⁷

Taken together, the analyses presented in this section suggest that our estimates of switcher elasticities are not biased, or at least not *upward* biased, by selection into job switching.

5.4 Are Earnings Elasticities Mediated by Firm-Specific Wage Premia?

Our approach to estimating earnings elasticities from job switchers uses variation from both firm and occupation transitions. As shown in Appendix Figure A.3, the earnings responses to lower taxes are similar for firm and occupation switchers.³⁸ In this section, we focus on firm switchers and ask if their earnings responses are mediated by firm-level wage effects as studied in the literature on AKM models (Abowd, Kramarz, and Margolis 1999). That is, while our quasi-experimental estimates should be interpreted as worker responses (as they are based on tax variation across workers, not firms), they may be mediated by job switchers sorting into higher-wage firms following the tax reform. This would be a different mechanism than the one modeled in section 2, albeit consistent with our general emphasis on the importance of job switching for earnings responses.

To investigate the role of firm-level effects, we estimate a standard AKM model of the form

$$\log z_{it} = \alpha_i + \psi_{J(i,t)} + \mathbf{X}_{it}\boldsymbol{\beta} + \nu_{it}, \quad (22)$$

where α_i is an individual fixed effect, $\psi_{J(i,t)}$ is a firm fixed effect, and \mathbf{X}_{it} is a vector of time-varying controls. The controls include year dummies, age dummies, and dummies for tenure in the individual's current firm. We estimate the model using pre-reform data (2002-2005), restricting the sample to firms with at least 10 employees. We merge the estimated firm coefficients $\hat{\psi}_{J(i,t)}$

³⁶This corresponds to the definition of mass layoffs in Jacobson, LaLonde, and Sullivan (1993). The definition is meaningful only for larger firms, so we further restrict the sample to firms with at least 20 employees at the time of the mass layoff. About one-fifth of job switchers in our baseline sample satisfy these mass-layoff criteria. Alternative definitions of mass layoffs give qualitatively similar results, but stricter definitions (increasing the minimum fraction laid off and/or the minimum number of employees) reduce sample size and increase standard errors.

³⁷The positive earnings effect of lower taxes among mass-layoff movers is consistent with a negative earnings effect of the mass layoff itself, as documented in the literature on job displacement. To confirm this, Appendix Figure A.5 provides an event study of mass layoffs, showing that they generate sizeable and persistent earnings losses. Hence, the findings in Figure 13 should be interpreted as showing that, within the group of switchers affected by mass layoffs, those who received tax cuts were less negatively affected than those who did not receive tax cuts.

³⁸Job switchers may change firm and occupation simultaneously. To retain statistical power, the subsamples in Figure A.3 include either all firm switches (even if occupation also changes) or all occupation switches (even if firm also changes).

onto our tax reform sample and run the following regression in the sample of firm switchers:

$$\Delta \hat{\psi}_{J(i,t)} = \beta_0 + \sum_{b \neq b_0} \beta_b \cdot \mathbf{1}[y_{i0} \in b] + \mu_{it}. \quad (23)$$

The difference-in-differences coefficient β_b captures the effect of the tax reform on the firm-specific earnings premia of firm switchers in income bin b . If the coefficients are positive in treated income bins, it implies that lower taxes induce switchers to sort into more remunerative firms, perhaps trading off non-wage amenities for higher wages.

The results are presented in Figure 14. It plots the changes in firm-specific earnings premia by income bin (the coefficients $\hat{\beta}_b$) in different time intervals: 2006-08 (placebo), 2008-10, 2010-12, and 2012-14. In every time interval, these changes are close to zero and statistically insignificant at all income levels above and below the treatment threshold. In other words, the earnings responses of firm switchers are not driven by tax-induced sorting across firms with different wage premia. This is consistent with our theoretical model in which earnings responses reflect dynamic returns to individual effort, realized at the point of switching.

6 Conclusion

The idea that the return to effort is dynamic seems *prima facie* true, especially for career workers at the top of the distribution. The very meaning of the word “career” contains a notion of dynamic progress. Yet, the issue of dynamic returns is largely ignored in the empirical literature on labor supply, presumably because of the challenges to estimating welfare-relevant, long-run elasticities in the presence of such returns. We have many compelling estimates of labor supply responses to tax reform, but they generally capture only short-term effects. In this paper, we propose a way to estimate long-term elasticities in the presence of dynamic returns without having to rely on a parameterized structural model.

We provide three contributions. First, we develop a new model of earnings responses to taxes in the presence of dynamic returns. In this model, the returns to effort are delayed and mediated by job switches such as promotions within firms or movements between firms. We use the model to provide a set of predictions that can be taken to the data, and to characterize how job switchers can be used to uncover the true long-run elasticity. Second, we provide descriptive evidence on earnings and hours-worked patterns over the lifecycle, verifying the predictions of the theoretical

model. This analysis leverages the granularity and statistical power of the Danish data to provide particularly clean evidence. Third, informed by the model and descriptive evidence, we conduct a quasi-experimental study of earnings elasticities using job switchers. A conventional estimation approach gives a modest earnings elasticity of about 0.2. Our job switcher approach, on the other hand, gives an elasticity in the range of 0.4-0.5. We present several analyses that address robustness and threats to identification, all of which support our empirical approach. A key advantage of the approach is that it does not require a unique experiment; it can be implemented using tax reform experiments of the type commonly used in the literature (as reviewed by [Saez, Slemrod, and Giertz 2012](#)).

While we argue that the long-run elasticity is larger than typically estimated, our analysis does not support the very large elasticities implied by some macro calibrations. What are the policy implications of an elasticity of 0.45 (our dynamic long-run estimate) as opposed to an elasticity of 0.2 (standard short-run estimate)? As an example, consider the Laffer rate on top earners. This is determined by the classic formula $\tau = 1 / (1 + \varepsilon a)$, where ε is the earnings elasticity and a is the Pareto parameter ([Diamond 1998](#); [Saez 2001](#)). Based on a Pareto parameter of 1.5 — the relevant number for the US — an elasticity of 0.2 implies $\tau = 0.77$, while an elasticity of 0.45 implies $\tau = 0.60$. Therefore, the long-run elasticity we estimate has significant policy implications.³⁹ The implications are even larger in countries with more compressed earnings distributions because their larger Pareto parameter magnifies the impact of the elasticity.

³⁹Still, even with our long-run elasticity, the current top marginal tax rate in the US is well below the Laffer point, which is also the optimal top tax rate absent any non-tax externalities.

References

- ABOWD, JOHN M., AND DAVID CARD (1987): "Intertemporal Labor Supply and Long-Term Employment Contracts," *The American Economic Review*, 77(1), 50–68. 6
- (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–445. 6
- ABOWD, JOHN M., FRANCIS KRAMARZ, AND DAVID N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251–333. 2, 6, 31
- ALTONJI, JOSEPH G., AND CHRISTINA H. PAXSON (1986): "Job Characteristics and Hours of Work," *Research in Labor Economics*, 8A, 1–55. 6
- (1988): "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs," *Journal of Labor Economics*, 6(2), 254–276. 6
- ANAGOL, SANTOSH, ALLAN DAVIDS, BENJAMIN B. LOCKWOOD, AND TARUN RAMADORAI (2022): "Diffuse Bunching with Frictions: Theory and Estimation," Working Paper. 6
- BEST, MICHAEL CARLOS, AND HENRIK KLEVEN (2013): "Optimal Income Taxation with Career Effects of Work Effort," Working Paper. 2
- BLUNDELL, RICHARD, AND THOMAS MACURDY (1999): "Labor Supply: A Review of Alternative Approaches," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3A, chap. 27. Elsevier. 4, 29
- BRONSON, MARY ANN, AND PETER SKOGMAN THOURSIE (2021): "The Wage Growth and Within-Firm Mobility of Men and Women: New Evidence and Theory," Working Paper. 3
- CABALLERO, RICARDO J., AND EDUARDO M.R.A. ENGEL (1991): "Dynamic (S,s) Economies," NBER Working Paper Working No. 3734. 8
- CALVO, GUILLERMO A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12(3), 383–398. 8
- CAPLIN, ANDREW, AND JOHN LEAHY (1991): "State-Dependent Pricing and the Dynamics of Money and Output," *The Quarterly Journal of Economics*, 106(3), 683–708. 8

- CAPLIN, ANDREW S., AND DANIEL F. SPULBER (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 102(4), 703–726. [8](#)
- CARD, DAVID, ANA RUTE CARDOSO, JOERG HEINING, AND PATRICK KLINE (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36(S1), S13–S70. [2, 6](#)
- CARD, DAVID, ANA RUTE CARDOSO, AND PATRICK KLINE (2016): "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, 131(2), 633–686. [2, 6](#)
- CARD, DAVID, JÖRG HEINING, AND PATRICK KLINE (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128(3), 967–1015. [2, 6](#)
- CHETTY, RAJ (2012): "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80(3), 969–1018. [5](#)
- CHETTY, RAJ, JOHN N. FRIEDMAN, TORE OLSEN, AND LUIGI PISTAFERRI (2011): "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *The Quarterly Journal of Economics*, 126(2), 749–804. [6](#)
- CHETTY, RAJ, JOHN N. FRIEDMAN, AND EMMANUEL SAEZ (2013): "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings," *American Economic Review*, 103(7), 2683–2721. [6](#)
- CHETTY, RAJ, ADAM GUREN, DAY MANOLI, AND ANDREA WEBER (2013): "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 27(1), 1–56. [5](#)
- DIAMOND, PETER A. (1998): "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *The American Economic Review*, 88(1), 83–95. [1, 7, 33](#)
- FARBER, HENRY S. (1999): "Mobility and Stability: The Dynamics of Job Change in Labor Markets," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3B, chap. 37. Elsevier. [6](#)
- GRIGSBY, JOHN, ERIK HURST, AND AHU YILDIRMAZ (2021): "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 111(2), 428–471. [2](#)

- GRUBER, JON, AND EMMANUEL SAEZ (2002): “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, 84(1), 1–32. [28](#)
- GUVENEN, FATIH, FATIH KARAHAN, SERDAR OZKAN, AND JAE SONG (2021): “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?,” *Econometrica*, 89(5), 2303–2339. [2](#)
- HARRIS, MILTON, AND BENGT HOLMSTRÖM (1982): “A Theory of Wage Dynamics,” *The Review of Economic Studies*, 49(3), 315–333. [6](#)
- HOLMSTRÖM, BENGT (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *The Review of Economic Studies*, 66(1), 169–182. [6](#)
- HUMLUM, ANDERS (2021): “Robot Adoption and Labor Market Dynamics,” Working Paper. [15](#)
- IMAI, SUSUMU, AND MICHAEL P. KEANE (2004): “Intertemporal Labor Supply and Human Capital Accumulation,” *International Economic Review*, 45(2), 601–641. [5](#)
- JACOBSON, LOUIS S., ROBERT J. LALONDE, AND DANIEL G. SULLIVAN (1993): “Earnings Losses of Displaced Workers,” *The American Economic Review*, 83(4), 685–709. [4](#), [30](#), [31](#)
- JAKOBSEN, KATRINE MARIE, AND JAKOB EGHOLT SØGAARD (2022): “Identifying Behavioral Responses to Tax Reforms: New Insights and a New Approach,” *Journal of Public Economics*, 212. [5](#), [22](#), [24](#)
- KEANE, MICHAEL, AND RICHARD ROGERSON (2015): “Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective,” *Annual Review of Economics*, 7(1), 89–117. [1](#), [2](#), [5](#)
- KEANE, MICHAEL P. (2011): “Labor Supply and Taxes: A Survey,” *Journal of Economic Literature*, 49(4), 961–1075. [2](#), [5](#)
- KLEVEN, HENRIK (2016): “Bunching,” *Annual Review of Economics*, 8, 435–464. [6](#)
- (2021): “Sufficient Statistics Revisited,” *Annual Review of Economics*, 13, 515–538. [10](#)
- (2023): “The EITC and the Extensive Margin: A Reappraisal,” NBER Working Paper No. 26405. [5](#)

- KLEVEN, HENRIK, MARTIN B. KNUDSEN, CLAUS THUSTRUP KREINER, SØREN PEDERSEN, AND EMMANUEL SAEZ (2011): “Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark,” *Econometrica*, 79, 651–692. [15](#)
- KLEVEN, HENRIK, CLAUS THUSTRUP KREINER, AND EMMANUEL SAEZ (2009): “The Optimal Income Taxation of Couples,” *Econometrica*, 77(2), 537–560. [7](#)
- KLEVEN, HENRIK, AND MAZHAR WASEEM (2013): “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan,” *The Quarterly Journal of Economics*, 128(2), 669–723. [6](#)
- KREINER, CLAUS THUSTRUP, SØREN LETH-PETERSEN, AND PEER EBBESEN SKOV (2016): “Tax Reforms and Intertemporal Shifting of Wage Income: Evidence from Danish Monthly Payroll Records,” *American Economic Journal: Economic Policy*, 8(3), 233–257. [5](#)
- KREINER, CLAUS THUSTRUP, JAKOB ROLAND MUNCH, AND HANS JØRGEN WHITTA-JACOBSEN (2015): “Taxation and the Long-Run Allocation of Labor: Theory and Danish Evidence,” *Journal of Public Economics*, 127, 74–86. [6](#)
- LABANCA, CLAUDIO, AND DARIO POZZOLI (2022): “Constraints on Hours within the Firm,” *Journal of Labor Economics*, 40(2), 473–503. [5](#), [6](#)
- LACHOWSKA, MARTA, ALEXANDRE MAS, RAFFAELE SAGGIO, AND STEPHEN A. WOODBURY (2023): “Work Hours Mismatch,” Working Paper. [6](#)
- NAKAMURA, EMI, AND JÓN STEINSSON (2010): “Monetary Non-Neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125(3), 961–1013. [8](#)
- PAVAN, RONNI (2011): “Career Choice and Wage Growth,” *Journal of Labor Economics*, 29(3), 549–587. [2](#), [6](#)
- PRESCOTT, EDWARD C. (2004): “Why Do Americans Work So Much More Than Europeans?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(1), 2–13. [1](#)
- ROGERSON, RICHARD, AND JOHANNA WALLENIS (2009): “Micro and macro elasticities in a life cycle model with taxes,” *Journal of Economic Theory*, 144(6), 2277–2292. [1](#)
- SAEZ, EMMANUEL (2001): “Using Elasticities to Derive Optimal Income Tax Rates,” *The Review of Economic Studies*, 68(1), 205–229. [1](#), [33](#)

——— (2010): “Do Taxpayers Bunch at Kink Points?,” *American Economic Journal: Economic Policy*, 2(3), 180–212. [7](#)

SAEZ, EMMANUEL, MANOS MATSAGANIS, AND PANOS TSAKLOGLOU (2012): “Earnings Determination and Taxes: Evidence From a Cohort-Based Payroll Tax Reform in Greece,” *The Quarterly Journal of Economics*, 127(1), 493–533. [5](#)

SAEZ, EMMANUEL, JOEL SLEMROD, AND SETH H. GIERTZ (2012): “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Literature*, 50(1), 3–50. [4](#), [10](#), [22](#), [24](#), [29](#), [33](#)

SCHEUER, FLORIAN, AND IVÁN WERNING (2017): “The Taxation of Superstars,” *The Quarterly Journal of Economics*, 132(1), 211–270. [6](#)

SIGAARD, HANS SCHYTTE (2022): “Labor Supply Responsiveness to Tax Reforms,” Working Paper. [5](#)

TAYLOR, JOHN B. (1980): “Output and Price Stability: An International Comparison,” *Journal of Economic Dynamics and Control*, 2, 109–132. [8](#)

TOPEL, ROBERT H., AND MICHAEL P. WARD (1992): “Job Mobility and the Careers of Young Men,” *The Quarterly Journal of Economics*, 107(2), 439–479. [2](#), [6](#)

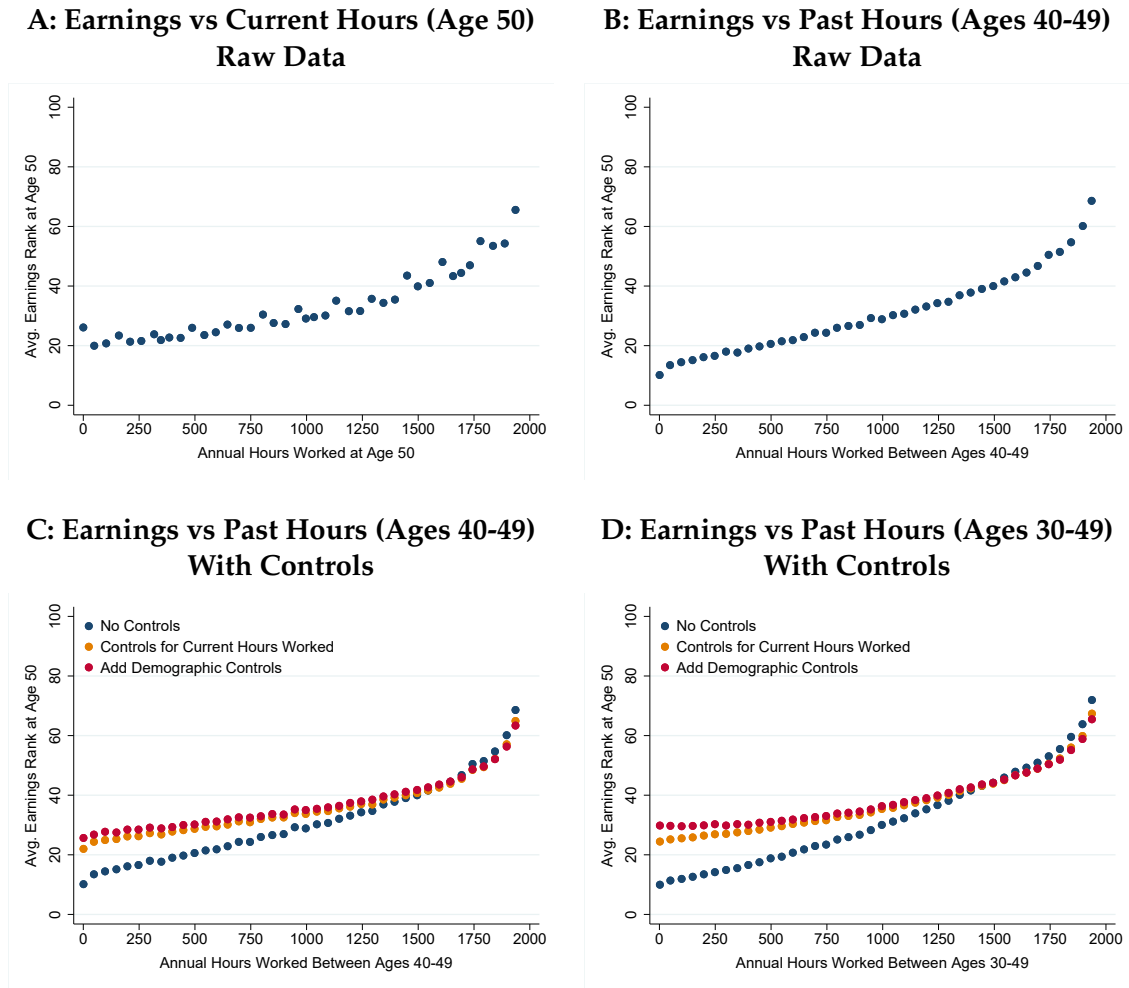
TORTAROLO, DARIO, GUILLERMO CRUCES, AND VICTORIA CASTILLO (2020): “It takes two to tango: Labor responses to an income tax holiday in Argentina,” Working Paper. [5](#)

TABLE 1: ELASTICITY ESTIMATES
DYNAMIC APPROACH VS STANDARD APPROACH

	Dynamic Elasticity		Standard Elasticity	
	ITT	TOT	ITT	TOT
<i>Time Period:</i>				
2008-10	0.208 (0.011)	0.412 (0.018)	0.086 (0.005)	0.171 (0.007)
2008-11	0.214 (0.011)	0.441 (0.019)	0.103 (0.005)	0.211 (0.009)
2008-12	0.217 (0.011)	0.468 (0.020)	0.115 (0.006)	0.248 (0.011)
2008-13	0.184 (0.011)	0.418 (0.022)	0.106 (0.006)	0.240 (0.013)
2008-14	0.176 (0.011)	0.412 (0.024)	0.109 (0.007)	0.255 (0.015)

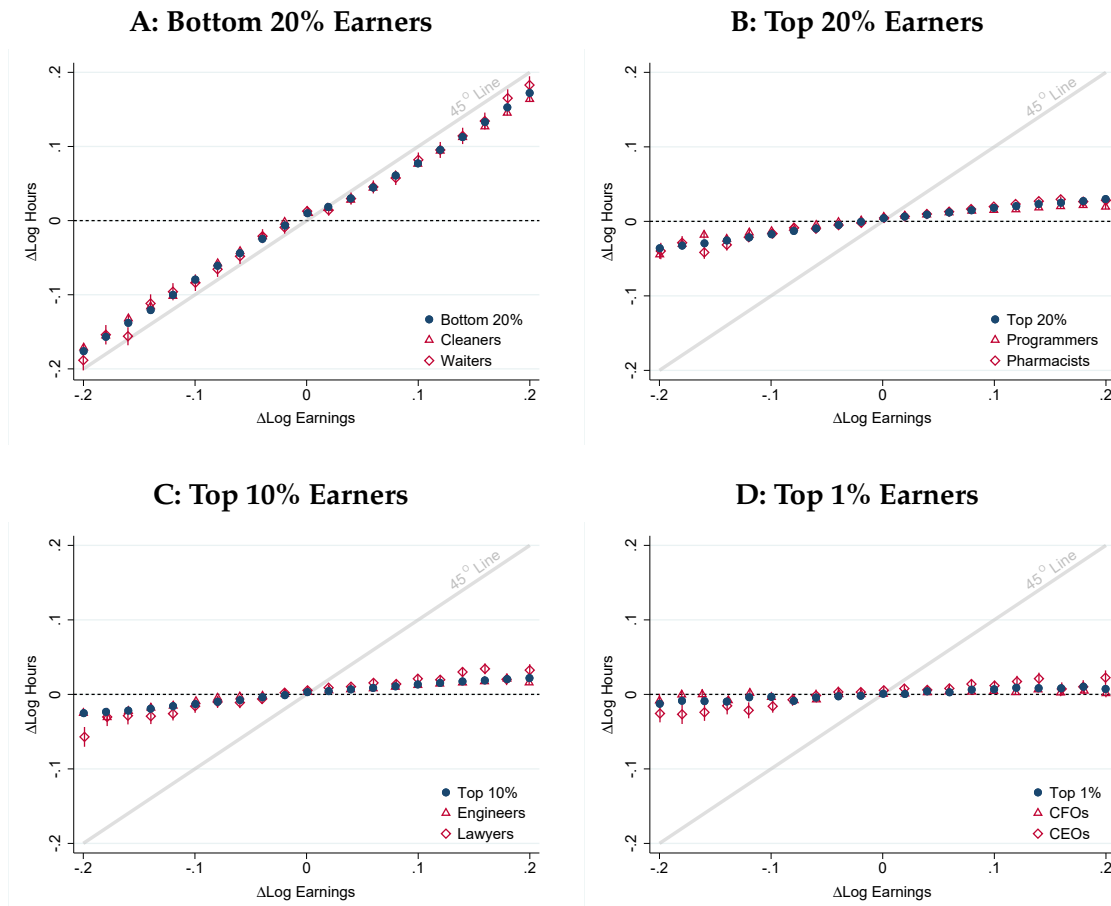
Notes: This table converts our estimates of earnings responses to the 2009 tax reform into elasticities with respect to the marginal net-of-tax rate. The table provides estimates based on the dynamic approach (left panel) and the standard approach (right panel) over time. The dynamic approach considers earnings responses among job movers (those who switch between firm \times occupation cells), while the standard approach considers earnings responses among all workers (movers and stayers). The relationship between the dynamic and standard elasticity estimates is characterized in equation (7). The table shows both intention-to-treat (ITT) and treatment-on-the-treated (TOT) estimates for each approach. The estimates of primary interest (highlighted in boldface) are the dynamic TOT estimates. These elasticities are large and stable over time, consistent with point identification of the long-run elasticity as shown in Proposition 4 and Corollary 1. Robust standard errors are provided in parentheses.

FIGURE 1: PAST HOURS WORKED PREDICT CURRENT EARNINGS, CONDITIONAL ON CURRENT HOURS WORKED



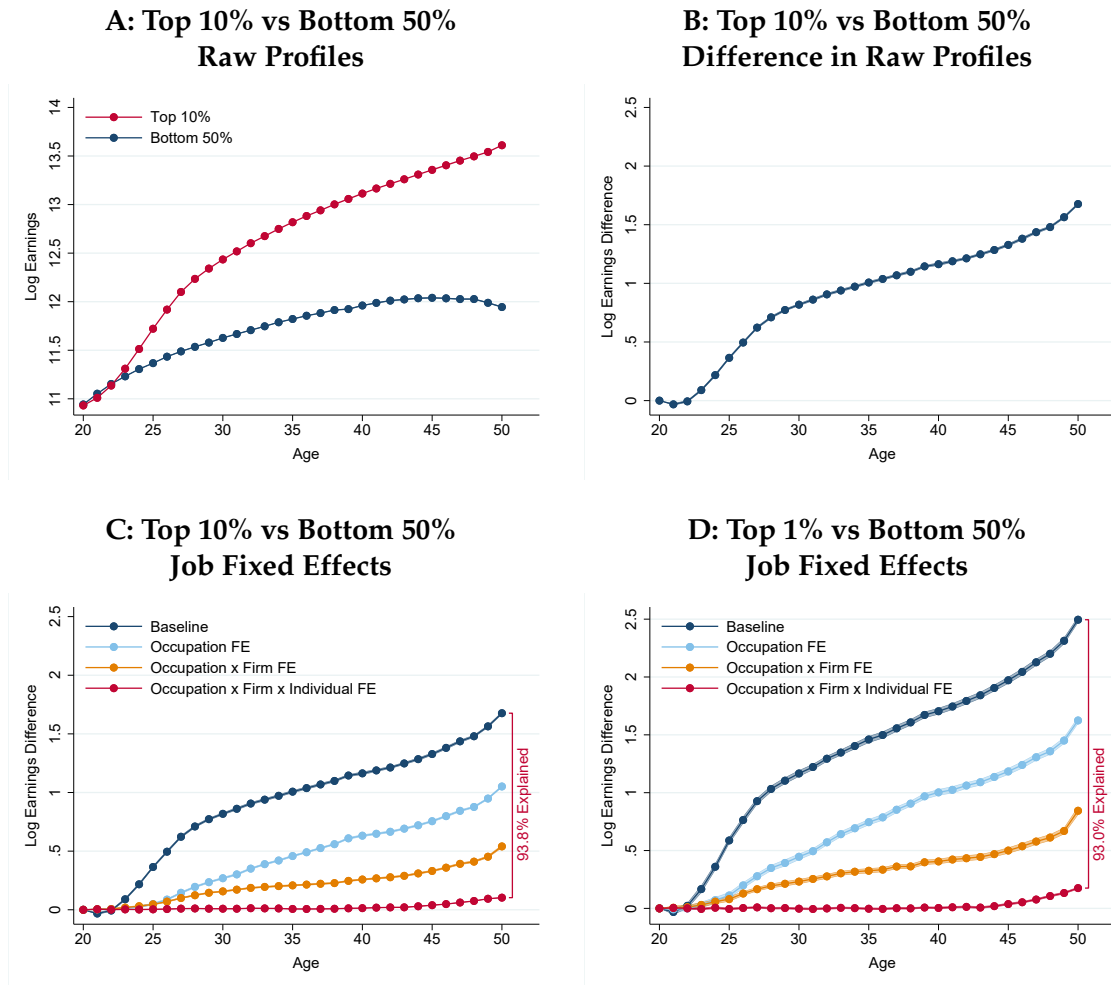
Notes: This figure shows that past hours worked predict current earnings. The figure is based on a balanced panel of workers observed between the ages of 20 and 50. Each panel plots the relationship between average earnings rank at age 50 and hours worked at different ages: current hours in Panel A and past hours in Panels B-D. The top panels depict raw data, while the bottom panels add controls. The predictive power of past hours worked remains strong even after controlling for current hours worked and demographic characteristics (dummies for gender, children, marital status, and education level). The graphs include 95% confidence intervals based on standard errors clustered at the individual level, but these are hardly visible.

FIGURE 2: CONTEMPORANEOUS HOURS AND EARNINGS CHANGES ARE UNRELATED AT THE TOP, BUT NOT AT THE BOTTOM




Notes: This figure shows the contemporaneous relationship between hours and earnings changes at the intensive margin. It plots changes in log hours against changes in log earnings in different segments of the earnings distribution: the bottom 20%, the top 20%, the top 10%, and the top 1%. The average relationship in each segment is depicted by blue dots, while examples of representative occupations in the different segments are depicted by red triangles and diamonds. While hours and earnings changes are almost perfectly correlated at the bottom (consistent with hourly-paid workers), they are virtually uncorrelated at the top (consistent with salaried workers). The error bars depict 95% confidence intervals based on standard errors clustered at the individual level.

FIGURE 3: LIFECYCLE PROFILES IN EARNINGS ARE DRIVEN BY DISCRETE JOB SWITCHES



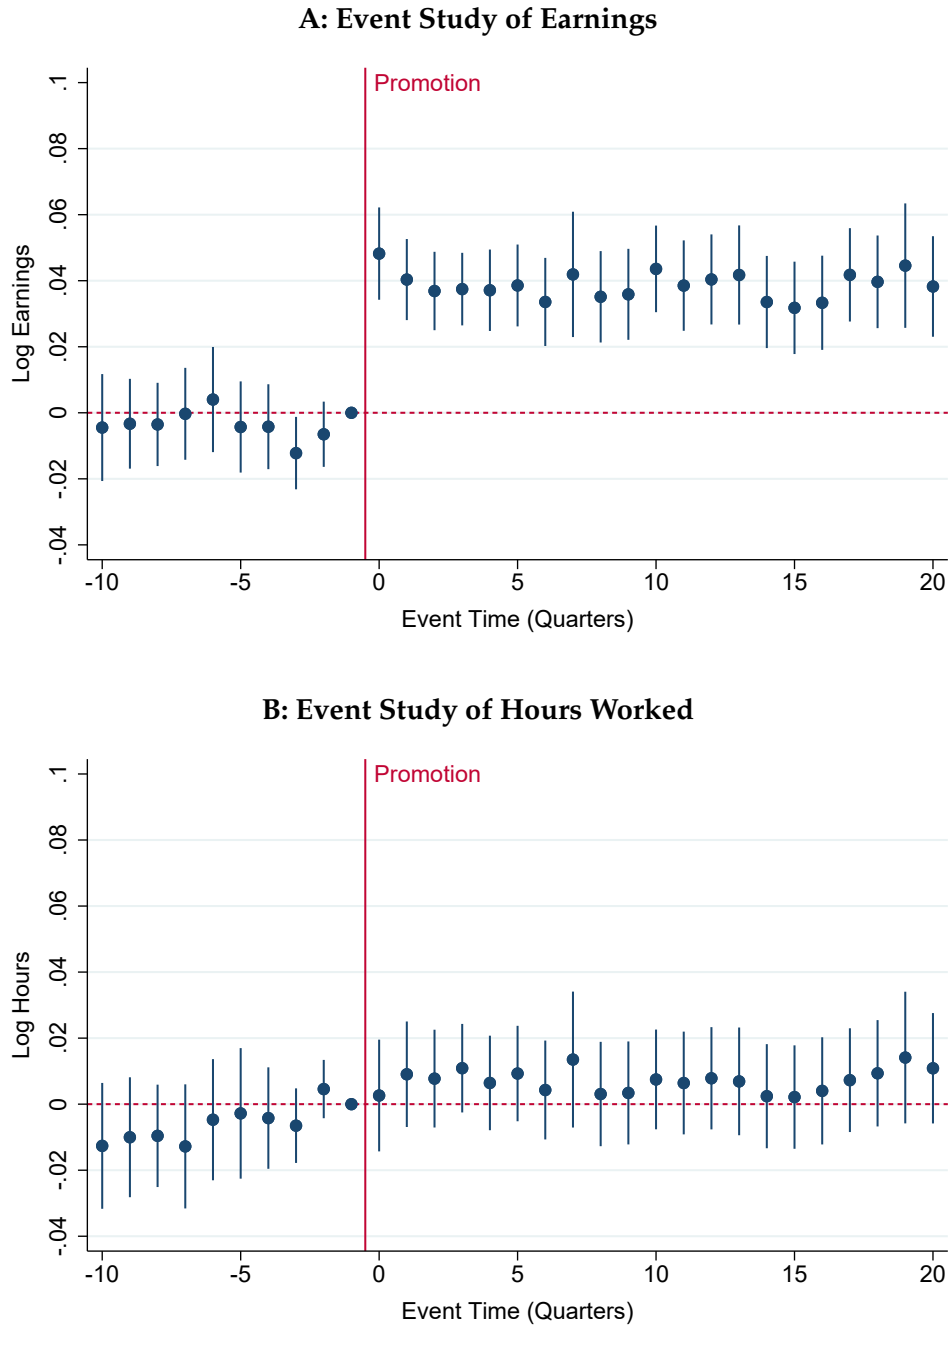
Notes: This figure shows that the lifecycle profile of earnings for workers who reach the top of the distribution is driven by job switches, defined as transitions between occupation×firm cells. The figure is based on a balanced panel of workers observed between the ages of 20 and 50, plotting the earnings profiles of workers observed in different quantiles of the distribution at age 50. Panel A plots the raw profiles of workers in the top 10% and bottom 50%, respectively. Panel B plots the difference in these raw profiles. Panel C compares the difference in raw profiles (dark blue) to the differences in net of occupation fixed effects (light blue), net of occupation×firm fixed effects (orange), and net of occupation×firm×individual fixed effects (red). Within individual, job fixed effects explain about 94% of the lifecycle divergence between top-10% and bottom-50% earners. Panel D repeats the analysis of Panel C, but for top-1% earners. The shaded areas (not always visible) represent 95% confidence intervals based on standard errors clustered at the individual level.

FIGURE 4: BETWEEN-JOB VS WITHIN-JOB VARIATION IN EARNINGS
 VARIANCE DECOMPOSITION BY EARNINGS RANK AT AGE 50




Notes: This figure decomposes the variance of lifecycle earnings into between-job variance and within-job variance using equation (18). As in Figure 3, jobs are measured as occupation×firm cells and the sample is a balanced panel of workers observed between the ages of 20 and 50. The figure plots the total variance of earnings (blue), the between-job variance (red), and the within-job variance (yellow) by earnings rank at age 50. At all ranks shown, almost all of the lifecycle variation in earnings can be attributed to between-job variation, i.e. to switches between occupation×firm cells. The shaded areas depict 95% confidence intervals based on standard errors clustered at the individual level.

FIGURE 5: EARNINGS JUMP DISCRETELY AT PROMOTIONS, WITH NO CHANGE IN HOURS




Notes: This figure presents event studies of promotions using specification (19). A promotion is defined as a switch to an occupation cell in which median earnings are at least 10% higher. The event study series show the outcomes of promoted workers relative to their unpromoted co-workers by quarter, normalizing the pre-promotion quarter to zero. Panel A considers earnings and Panel B considers hours worked. Promotions lead to sharp jumps in earnings, while hours worked are smooth. The error bars depict 95% confidence intervals based on standard errors clustered at the individual level.

FIGURE 6: 2009 TAX REFORM IN DENMARK



Notes: This figure illustrates the tax variation created by the 2009 reform in Denmark. The reform reduced marginal tax rates above an income threshold by 11.4pp on average (as depicted by the red series), while leaving marginal tax rates below the income threshold virtually unchanged (as depicted by the black series). The threshold that separates treatments and controls is located at around the 70th percentile of the income distribution.

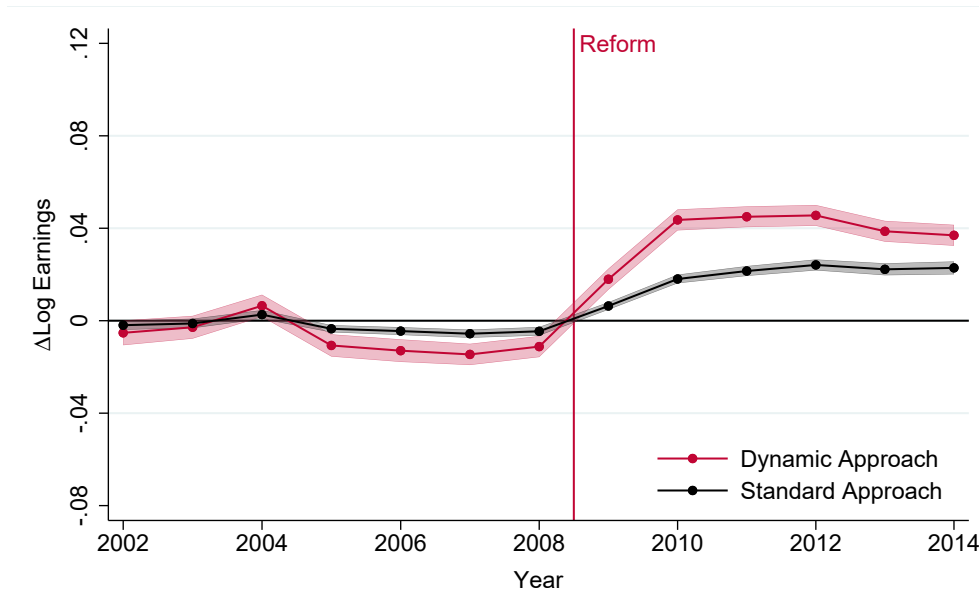
FIGURE 7: IMPACT OF TAX REFORM ON EARNINGS




Notes: This figure shows the impact of the 2009 tax reform on earnings for all workers (Panel A) and for job movers vs job stayers (Panel B). Job movers are defined as those who switch between firm \times occupation cells. Each panel plots changes in log earnings between 2008-10 (actual experiment) and between 2006-08 (placebo experiment) by pre-reform income bin, depicting the threshold that separates treatments and controls by a vertical line. The 2006-08 series capture non-tax effects on earnings growth that vary by income level (including mean reversion). The 2008-10 and 2006-08 series track each other below the treatment threshold, but diverge above the threshold. This divergence represents the causal impact of the tax reform on earnings. Panel A shows a precisely estimated, but modest earnings increase in the full sample of treated workers. Panel B shows that the modest average impact masks strong heterogeneity by job switching status. Job movers feature large earnings increases, while job stayers feature no earnings increases. The shaded areas show 95% confidence intervals based on robust standard errors.

FIGURE 8: IMPACT OF TAX REFORM ON EARNINGS OVER TIME
 DYNAMIC APPROACH (SWITCHERS) VS STANDARD APPROACH (ALL WORKERS)

A: Intention-to-Treat (ITT)




B: Treatment-on-the-Treated (TOT)



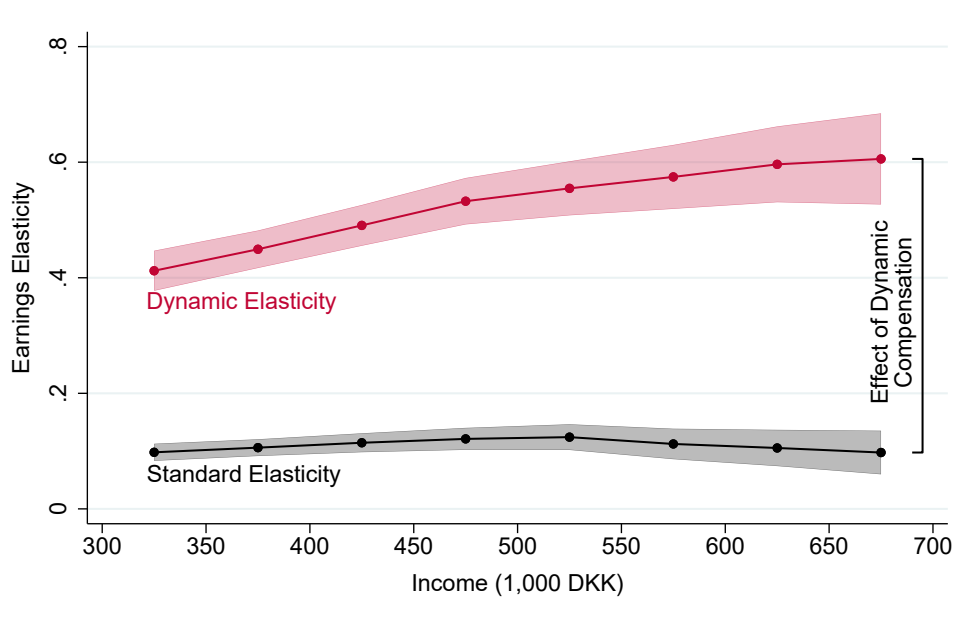
Notes: This figure shows the time profile of ITT effects (Panel A) and TOT effects (Panel B) on earnings. The red series show earnings responses among job movers (“dynamic approach”). These are based on quadruple-differences estimates that compare earnings growth above vs below the treatment threshold for job movers vs job stayers in a post-period relative to a pre-period. The pre-reform part of the series ($t = 2002, \dots, 2008$) gives placebo estimates, while the post-reform part of the series ($t = 2009, \dots, 2014$) gives treatment impact estimates among workers who have made their first post-reform job move by year t . The post-reform mover estimates are stable over time, consistent with point identification of the long-run earnings response as shown in Proposition 4 and Corollary 1. The black series show earnings responses in the full sample of workers, including job stayers (“standard approach”). These are simulated by scaling the estimates for job movers using the cumulated fraction of workers who have made a post-reform job switch, as characterized by equation (7). The earnings responses in the full sample increase over time as more workers make job transitions, but the convergence between the standard and dynamic approaches is slow. The shaded areas depict 95% confidence intervals based on robust standard errors.

FIGURE 9: IMPACT OF TAX REFORM ON SWITCHING PROBABILITY



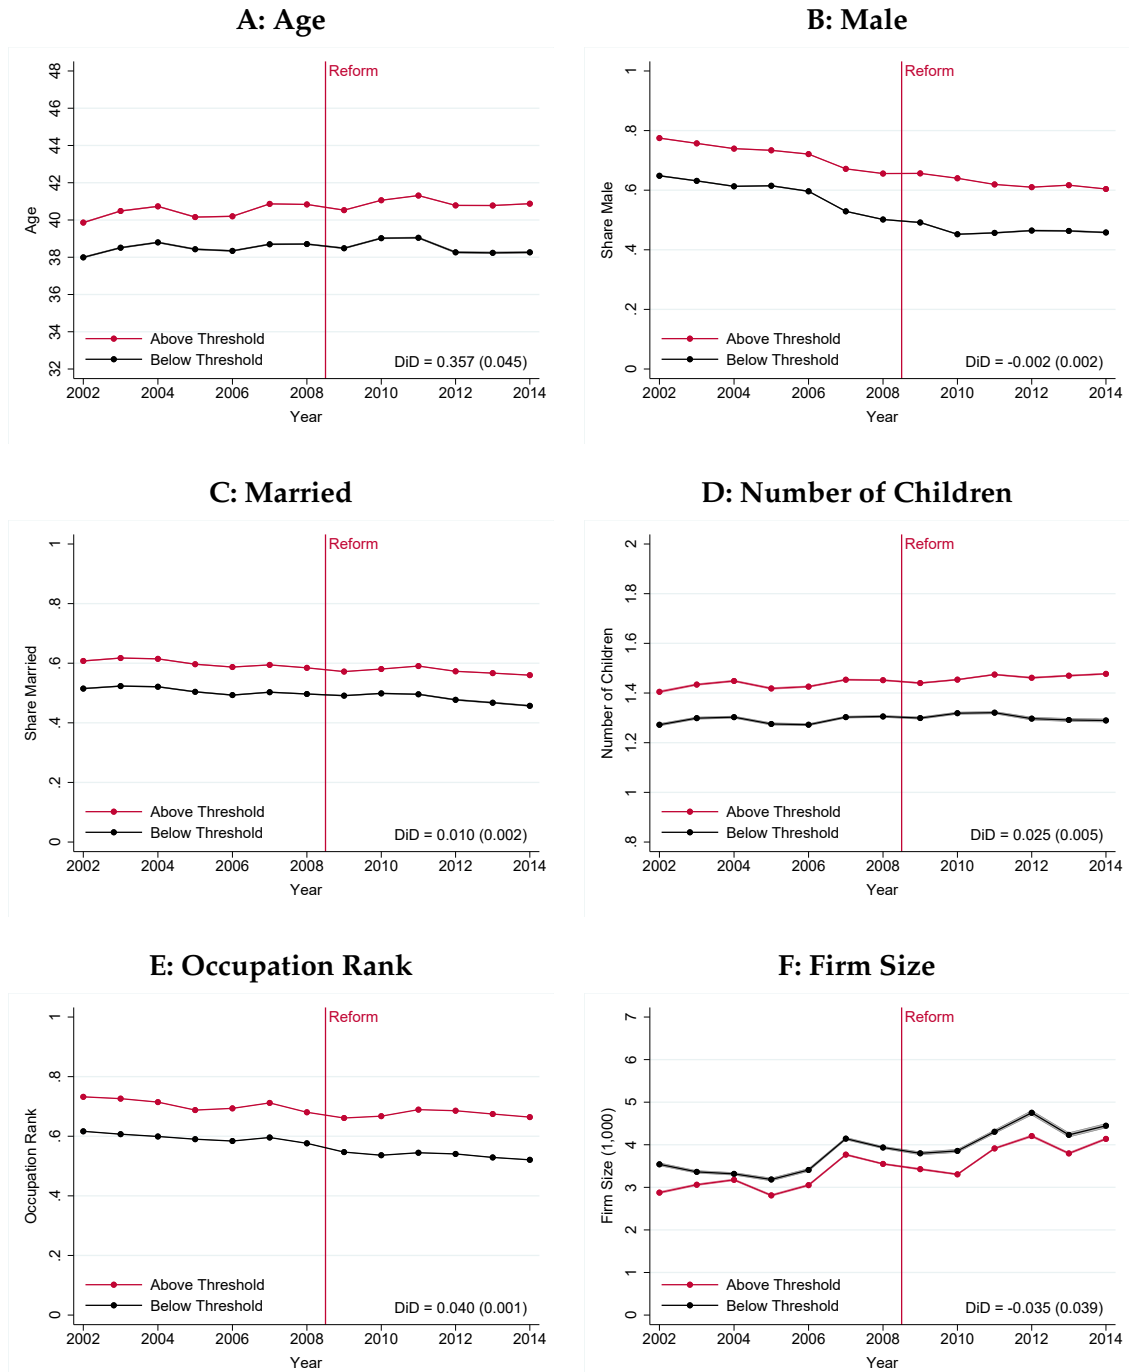
Notes: This figure investigates if the 2009 tax reform impacted the probability of switching jobs, defined as moving between firm×occupation cells. It plots switching probabilities for workers above and below the treatment threshold between 2002 and 2014, depicting the reform by a vertical line. The switching probability for workers below the threshold (controls) has been normalized to the level for workers above the threshold (treatments) in the pre-reform year, 2008. The figure also reports a difference-in-differences estimate of the average effect on the switching probabilities of treated workers over the post-reform period. This effect is extremely small and precisely estimated. The shaded areas (hardly visible) show 95% confidence intervals based on robust standard errors.

FIGURE 10: DYNAMIC VS STANDARD ELASTICITIES BY INCOME LEVEL



Notes: This figure shows earnings elasticities by income level based on the standard approach (black series) and the dynamic approach (red series). The standard approach considers the average effect on all workers, while the dynamic approach considers the effect on job movers only. For each approach, the figure shows elasticities between 2008-10 above different income cutoffs, depicted on the x-axis. The standard elasticity is small and virtually constant in income. The dynamic elasticity is large and increasing in income. The gap between the two — the effect of dynamic compensation realized at the point of switching — is therefore increasing in the level of income. The shaded areas show 95% confidence intervals based on robust standard errors.


FIGURE 11: IS SWITCHING SELECTED?
 IMPACT OF TAX REFORM ON SWITCHER CHARACTERISTICS




Notes: This figure investigates if the observable characteristics of treated job switchers relative to untreated job switchers diverge after the 2009 tax reform. Each panel plots the time series of a demographic variable for job switchers above and below the treatment threshold. Six different variables are considered: age, fraction male, fraction married, number of children, occupational rank (based on average earnings within each occupation cell), and firm size. The figure also reports difference-in-differences estimates of the effects on the different demographic variables. These effects are very small, albeit statistically significant in several cases due to the statistical power of our data. The analysis shows that job switching is not selected on observables given our quasi-experimental design. The shaded areas (hardly visible) show 95% confidence intervals based on robust standard errors.

FIGURE 12: IMPACT OF TAX REFORM ON EARNINGS OVER TIME
ROBUSTNESS OF DYNAMIC APPROACH TO CONTROLS

A: Intention-to-Treat (ITT)

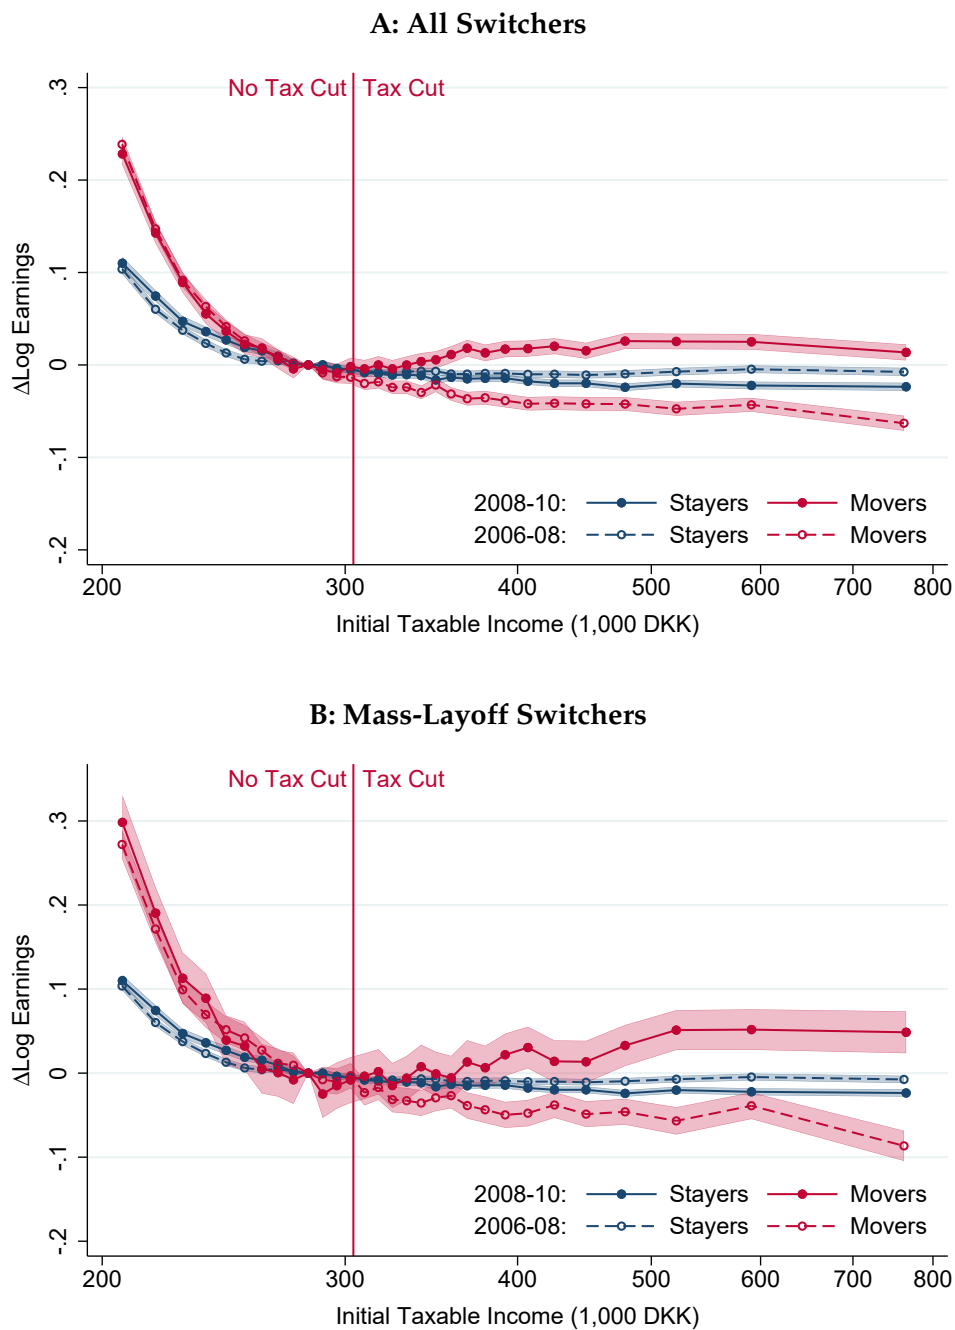


B: Treatment-on-the-Treated (TOT)




Notes: This figure investigates the robustness of the results in Figure 8 to control variables. It plots the time profile of earnings effects (ITT and TOT, respectively) among job switchers under different sets of controls. Four different specifications are considered: no controls (dark blue), demographic controls (light blue), fixed effects for initial job cell (orange), and finally job cell fixed effects and demographic controls combined (red). The demographic variables include dummies for age, gender, marital status, and number of children. The specification with only job cell fixed effects corresponds to our baseline results in Figure 8. The analysis shows that our dynamic approach is robust to the specification of controls, consistent with the finding in the preceding figure that job switching is not selected on observables. The error bars depict 95% confidence intervals based on robust standard errors.

FIGURE 13: IS SWITCHING SELECTED?
ALL SWITCHERS VS MASS-LAYOFF SWITCHERS



Notes: This figure shows the impact of the 2009 tax reform on earnings in the full sample of switchers (Panel A) and in a sample of mass-layoff switchers (Panel B). Switchers are those who move firm and/or occupation, while mass-layoff switchers are those who move firm following a mass layoff in their original firm. To qualify as a mass layoff, we require that a firm reduces its workforce by at least 30% and has at least 20 employees ex ante. Each panel plots changes in log earnings between 2008-10 (actual experiment) and between 2006-08 (placebo experiment) by pre-reform income bin, depicting the threshold that separates treatments and controls by a vertical line. The earnings responses in the mass-layoff sample are similar to those in the full sample of switchers. This suggests that our estimates of earnings elasticities among job switchers are not biased by selection on unobservables. The shaded areas show 95% confidence intervals based on robust standard errors.

FIGURE 14: IMPACT OF TAX REFORM ON FIRM-LEVEL WAGE PREMIA FOR SWITCHERS



Notes: This figure investigates if the earnings responses to lower taxes are mediated by firm-level wage premia. The estimation of firm-level wage premia is based on the AKM model in equation (22). Using the estimated wage premia as outcomes, the figure plots coefficients from the difference-in-differences specification (23) run on the sample of firm switchers. These coefficients represent changes in the firm-specific wage premia of firm switchers by income bin over different time intervals: 2006-08 (placebo), 2008-10, 2010-12, and 2012-14. In all series, the changes in firm-specific wage premia are close to zero and statistically insignificant in all income bins (relative to the omitted bin below the treatment threshold). Hence, the earnings responses of firm switchers are not driven by tax-induced sorting into firms with higher wage premia. The shaded areas show 95% confidence intervals based on robust standard errors.

Online Appendix

A Supplementary Figures and Tables


TABLE A.1: TOP VS BOTTOM OCCUPATIONS
JOB TITLES AND EARNINGS

Occupation	Earnings (DKK 1,000)			
	Mean	P10	P90	
Top 10	Top Executives	2,107.6	447.7	5,181.2
	Managing Directors	1,071.7	407.8	1,890.2
	Securities and Currency Traders	1,050.7	397.5	1,963.2
	Administrative Directors	1,024.0	382.2	1,803.6
	Lawyers	983.6	401.5	1,810.8
	Pilots	929.7	476.7	1,400.5
	Medical Doctors	899.9	486.3	1,287.4
	Senior Government Officials	871.4	492.8	1,432.2
	Finance and Insurance Analysts	849.7	442.8	1,332.6
	Managers, Police and Judiciary	843.6	703.2	1,041.3
Bottom 10	Retail Assistants	268.8	56.8	484.0
	Machine Operators	268.8	59.6	527.3
	Cleaners	266.0	180.3	355.6
	Street and Market Sales Persons	260.3	29.6	481.6
	Services and Sales Workers	258.1	73.7	460.8
	Tailors	257.2	57.6	439.3
	Couriers	256.1	122.2	399.2
	Pottery Makers	240.4	67.1	400.6
	Beauticians	235.6	57.5	438.5
	Manual Laborers, Agriculture	221.4	103.0	343.3


Notes: This table shows the highest-paying and lowest-paying occupations among workers aged 45-50. The classification is based on 6-digit occupation codes, ranked by mean wage earnings. For each occupation cell, the table reports the mean, the 10th percentile (P10), and the 90th percentile (P90) of wage earnings.

FIGURE A.1: VALIDATION OF ADMINISTRATIVE HOURS WORKED MEASURE
ADMINISTRATIVE DATA VS SURVEY DATA

A: Administrative Data (Pension Measure of Hours Worked)




B: Survey Data (Self-Reported Measure of Hours Worked)




Notes: This figure validates the administrative measure of hours worked — the pension measure described in section 3 — against a survey measure of hours worked. The survey measure is based on a question about actual, uncapped hours taken from the Danish component of the EU Labour Force Survey. The figure plots the relationship between earnings and hours worked in the administrative data (Panel A) and in the survey data (Panel B). The earnings-hours relationship is similar in the two data sources. However, the survey measure is much more noisy than the administrative measure, especially at the top of the hours and earnings distribution, which is a key reason for using administrative data. The error bars depict 95% confidence intervals.

FIGURE A.2: DISTRIBUTION OF NUMBER OF SWITCHES
TOP VS BOTTOM EARNERS BETWEEN AGES 20-50

A: Top 10% vs Bottom 50%



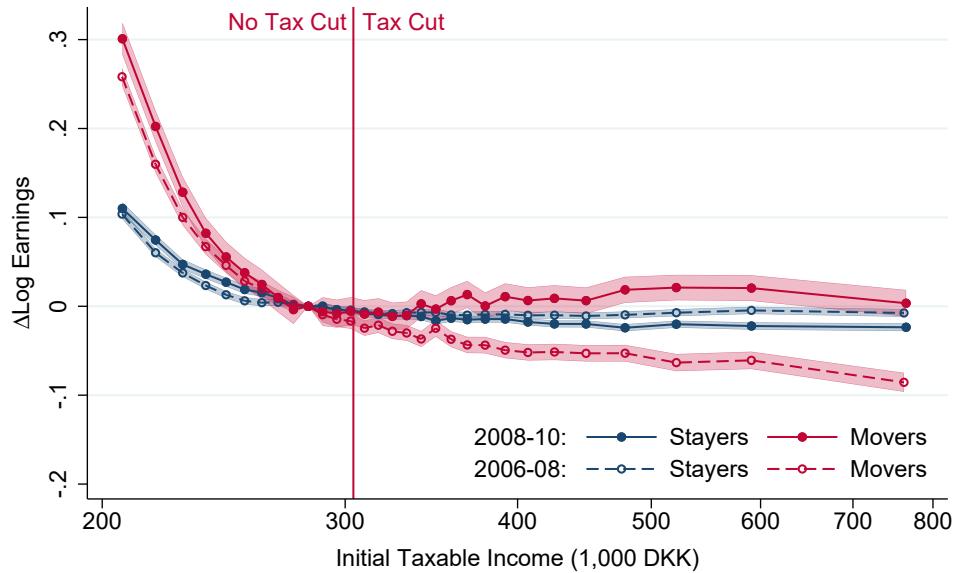
B: Top 1% vs Bottom 50%




Notes: This figure shows the distribution of the number of job switches between ages 20-50 for top and bottom earners. The figure is based on a balanced panel of workers observed between ages 20-50, splitting the sample by their earnings percentile at age 50. Panel A compares top-10% and bottom-50% earners, while Panel B compares top-1% and bottom-50% earners. The distributions are broadly similar across top and bottom earners. The average number of job switches is about 10 at the top and 9 at the bottom, corresponding to roughly one switch every three years.

FIGURE A.3: IMPACT OF TAX REFORM ON EARNINGS
 BY TYPE OF SWITCH

A: Firm Switches




B: Occupation Switches




Notes: This figure shows the impact of the 2009 tax reform on earnings for firm switchers (Panel A) and occupation switchers (Panel B), each of them compared to non-switchers. To retain statistical power, Panel A includes all firm switchers (even if they also switch occupation) while Panel B includes all occupation switchers (even if they also switch firm). The figure is otherwise constructed in the same way as Figure 7. It plots changes in log earnings between 2008-10 (actual experiment) and between 2006-08 (placebo experiment) by pre-reform income bin. The empirical patterns are very similar in the two samples, with large earnings responses to lower taxes among both firm and occupation switchers. The shaded areas show 95% confidence intervals based on robust standard errors.

**FIGURE A.4: IMPACT OF TAX REFORM ON SWITCHING PROBABILITY
BY TYPE OF SWITCH**

A: Firm Switches




B: Occupation Switches



Notes: This figure shows the impact of the 2009 tax reform on the probability of switching between firms (Panel A) and between occupations (Panel B). The figure is constructed in the same way as Figure 9 in which all switchers were pooled. It plots, for firm and occupation switches respectively, the switching probability for workers above and below the treatment threshold between 2002 and 2014. The switching probability for workers below the threshold (controls) has been normalized to the level for workers above the threshold (treatments) in the pre-reform year, 2008. The figure reports difference-in-differences estimates of the average effects on firm- and occupation-switching probabilities for treated workers over the post-reform period. These effects are extremely small and precisely estimated. The shaded areas (hardly visible) show 95% confidence intervals based on robust standard errors.

FIGURE A.5: IMPACT OF MASS LAYOFF ON EARNINGS



Notes: This figure presents an event study of the effect of mass layoffs on earnings. Mass layoffs are defined as layoffs in which firms with at least 20 employees reduce their workforce by at least 30% in a single year. The figure shows log earnings by event time (blue series) compared to a linear time trend estimated on pre-layoff data (dashed line). Mass layoffs lead to sizeable and persistent earnings losses. The shaded area depicts 95% confidence intervals based on standard errors clustered at the individual level.

B Theoretical Proofs

B.1 Proof of Proposition 1

We insert flow utility (1) into the objective (3), which gives the following maximization problem:

$$\max_{y_t} \left\{ \sum_{s=t}^{\infty} \delta^{s-t} \mathbb{E} [(1-\tau) z_s] - n_t v(y_t/n_t) \right\}.$$

The first-order condition with respect to y_t is given by

$$(1-\tau) \sum_{s=t}^{\infty} \delta^{s-t} \frac{d}{dy_t} \mathbb{E} [z_s] = v'(y_t/n_t). \quad (24)$$

Using equation (2) to substitute for $\mathbb{E} [z_s]$, we obtain

$$\lambda(1-\tau) \sum_{s=t}^{\infty} \delta^{s-t} (1-\lambda)^{s-t} = v'(y_t/n_t).$$

Given the parameterization $v(x) = \frac{\eta}{\eta+1} x^{\frac{\eta+1}{\eta}}$, this may be rewritten as

$$\lambda(1-\tau) \sum_{s=t}^{\infty} (\delta(1-\lambda))^{s-t} = (y_t/n_t)^{\frac{1}{\eta}}.$$

Finally, by using the relationship $\sum_{s=t}^{\infty} x^{s-t} = \frac{1}{1-x}$, we obtain the result in equation (4).

B.2 Proof of Proposition 2, Part 2

The correlation coefficient between z_t and y_t equals

$$\text{corr}(z_t, y_t) = \frac{\text{cov}(z_t, y_t)}{\sigma_{z_t} \sigma_{y_t}}, \quad (25)$$

where the covariance is defined as $\text{cov}(z_t, y_t) \equiv \mathbb{E} [(z_t - \bar{z}_t)(y_t - \bar{y}_t)]$. Using equation (2), this covariance may be written as

$$\begin{aligned}
\text{cov}(z_t, y_t) &= \mathbb{E}[(\lambda(y_t - \bar{y}_t) + (1 - \lambda)(z_{t-1} - \bar{z}_{t-1}))(y_t - \bar{y}_t)] \\
&= \mathbb{E}\left[\lambda(y_t - \bar{y}_t)^2 + (1 - \lambda)(z_{t-1} - \bar{z}_{t-1})(y_t - \bar{y}_t)\right] \\
&= \lambda \text{var}(y_t) + (1 - \lambda) \text{cov}(z_{t-1}, y_t) \\
&= \lambda \text{var}(y_t),
\end{aligned}$$

where we used that $\text{cov}(z_{t-1}, y_t) = 0$, because y_t depends only on the current realization of n_t , while z_{t-1} only depends on realizations of n_s for periods $s < t$.

To compute the correlation coefficient, we also use that $\sigma_{z_t}^2 = \lambda \sigma_{y_t}^2 + (1 - \lambda) \sigma_{z_{t-1}}^2$ from the earnings specification (2). This implies that $\sigma_{z_t}^2 = \lambda \sum_{s=0}^{\infty} (1 - \lambda)^s \sigma_{y_{t-s}}^2$. From equation (4) and $n_t = f(t) + \mu$, it follows that $\sigma_{y_t}^2$ is time-invariant, i.e. $\sigma_{y_t}^2 = \sigma_y^2$ for $\forall t$. Using this time-invariance along with the property $\sum_{s=0}^{\infty} x^s = \frac{1}{1-x}$, it follows that $\sigma_{z_t}^2 = \sigma_y^2$ and, hence, $\sigma_{z_t} = \sigma_y$. By inserting this property and the above formula for the covariance into the definition in (25), we obtain

$$\text{corr}(z_t, y_t) = \frac{\text{cov}(z_t, y_t)}{\sigma_{z_t} \sigma_{y_t}} = \frac{\lambda \sigma_{y_t}^2}{\sigma_{y_t}^2} = \lambda.$$

B.3 Social Welfare = Steady State Welfare When the Social Discount Factor is 1

Consider a social planner who wants to minimize the present discounted value of the deadweight loss from taxation, $\sum_{t=0}^{\infty} \rho^t D_t$, where ρ is the social discount factor. This objective is not well-defined for $\rho = 1$ and, therefore, we redefine the planner's objective function as

$$\Psi \equiv (1 - \rho) \sum_{t=0}^{\infty} \rho^t D_t. \quad (26)$$

Because this objective function is just a monotone transformation of the original objective, they will yield identical optimal solutions. By adding and subtracting the steady state value D^* , the objective may be rewritten as

$$\begin{aligned}
\Psi &= D^* + (1 - \rho) \sum_{t=0}^{\infty} \rho^t (D_t - D^*) \\
&= D^* + (1 - \rho) \sum_{t=0}^{T-1} \rho^t (D_t - D^*) + (1 - \rho) \sum_{t=T}^{\infty} \rho^t (D_t - D^*)
\end{aligned} \quad (27)$$

Given D_t is converging gradually towards D^* , the last term can be bounded:

$$\left| (1 - \rho) \sum_{t=T}^{\infty} \rho^t (D_t - D^*) \right| \leq |D_T - D^*| (1 - \rho) \sum_{t=T}^{\infty} \rho^t = |D_T - D^*| \rho^T.$$

By substituting this into equation (27), we obtain

$$\left| (1 - \rho) \sum_{t=0}^{\infty} \rho^t D_t - D^* \right| \leq (1 - \rho) \sum_{t=0}^{T-1} \rho^t |D_t - D^*| + |D_T - D^*| \rho^T \quad \forall T.$$

This implies

$$\begin{aligned} \lim_{\rho \rightarrow 1} \left| (1 - \rho) \sum_{t=0}^{\infty} \rho^t D_t - D^* \right| &\leq \lim_{\rho \rightarrow 1} (1 - \rho) \sum_{t=0}^{T-1} \rho^t (D_t - D^*) + \lim_{\rho \rightarrow 1} |D_T - D^*| \rho^T \quad \forall T \\ &\Leftrightarrow \lim_{\rho \rightarrow 1} \left| (1 - \rho) \sum_{t=0}^{\infty} \rho^t D_t - D^* \right| \leq |D_T - D^*| \quad \forall T. \end{aligned}$$

Because D_T converges to D^* as T increases, it follows that

$$\lim_{\rho \rightarrow 1} (1 - \rho) \sum_{t=0}^{\infty} \rho^t D_t = D^*.$$

Therefore, at a social discount factor of $\rho = 1$, the welfare objective in equation (26) is equivalent to steady state welfare D^* . In this case, welfare analysis and policy design depend only on steady state elasticities, not the contemporaneous elasticities typically estimated.

B.4 Generalization of Proposition 3

When deriving equation (7), we disregarded any systematic lifecycle trend in earnings, i.e., $g(t)$ was assumed to be constant. In the general case where we impose only the initial condition $\bar{y}_0 = z_{-1}$, we obtain from equation (5):

$$\varepsilon_t^z = \frac{\lambda \sum_{s=0}^t (1 - \lambda)^s \bar{y}_{t-s} \frac{d\bar{y}_{t-s}/\bar{y}_{t-s}}{d(1-\tau)/(1-\tau)}}{\lambda \sum_{s=0}^t (1 - \lambda)^s \bar{y}_{t-s} + (1 - \lambda \sum_{s=0}^t (1 - \lambda)^s) \bar{z}_{-1}}.$$

From equation (4), we have $\frac{d\bar{y}_{t-s}/\bar{y}_{t-s}}{d(1-\tau)/(1-\tau)} = \eta$. Hence,

$$\varepsilon_t^z = \alpha_t \eta,$$

where

$$\alpha_t = \frac{\lambda \sum_{s=0}^t (1-\lambda)^s \bar{y}_{t-s}}{\lambda \sum_{s=0}^t (1-\lambda)^s \bar{y}_{t-s} + (1-\lambda \sum_{s=0}^t (1-\lambda)^s) z_{-1}}.$$

In this general expression, it remains the case that α_t increases over time from $\alpha_0 = \lambda$ to $\alpha_\infty = 1$.

B.5 Proof of Proposition 5

The expected profits of hiring a worker at time t on a fixed-wage contract \hat{z}_t equals

$$\mathbb{E}[\pi] = \sum_{s=t}^{\infty} \delta^{s-t} \mathbb{E}[y_s - \hat{z}_t] (1-\lambda)^{s-t},$$

where $(1-\lambda)^{s-t}$ is the probability of retaining the worker in the same contract until time s . We assume free entry/exit of firms and that firms can pool risk. This implies that expected profits are zero in equilibrium. From the previous equation, we can solve for the competitive wage level as a function of expected worker output:

$$\hat{z}_t = \frac{\sum_{s=t}^{\infty} \delta^{s-t} \mathbb{E}[y_s] (1-\lambda)^{s-t}}{\sum_{s=t}^{\infty} \delta^{s-t} (1-\lambda)^{s-t}}. \quad (28)$$

The firm only observes worker output y_s at time $s = t$. Therefore, expected future output $\mathbb{E}[y_s]$ must align with the expected optimal choice of the worker. The first-order condition with respect to y_t is still given by equation (24). Using equation (13) to substitute for $\mathbb{E}[z_s]$ in equation (24) and noting that the wage will be fixed until the next job event occurs, we obtain

$$\lambda(1-\tau) \frac{d\mathbb{E}[\hat{z}_t]}{dy_t} \sum_{s=t}^{\infty} \delta^{s-t} (1-\lambda)^{s-t} = v'(y_t/n_t). \quad (29)$$

Since the productivity of the worker evolves according to $n_s = n_t(1+g)^{s-t}$, we guess that the solution is characterized by $y_s = y_t(1+g)^{s-t}$. We use this property to find a solution for (y_t, \hat{z}_t) and then verify that the property is in fact satisfied for this solution. Using the property, we may write the wage equation (28) as

$$\hat{z}_t = y_t \frac{\sum_{s=t}^{\infty} \delta^{s-t} (1+g)^{s-t} (1-\lambda)^{s-t}}{\sum_{s=t}^{\infty} \delta^{s-t} (1-\lambda)^{s-t}} = y_t \frac{1-\delta(1-\lambda)}{1-\delta(1+g)(1-\lambda)}.$$

This is equation (15) in Proposition 5. From this equation, we also get

$$\frac{d\hat{z}_t}{dy_t} = \frac{1 - \delta(1 - \lambda)}{1 - \delta(1 + g)(1 - \lambda)}.$$

By inserting this expression into the first-order condition (29), we obtain

$$\lambda(1 - \tau) \frac{1 - \delta(1 - \lambda)}{1 - \delta(1 + g)(1 - \lambda)} \sum_{s=t}^{\infty} \delta^{s-t} (1 - \lambda)^{s-t} = v'(y_t/n_t),$$

which gives

$$\frac{\lambda}{1 - \delta(1 - \lambda)(1 + g)} (1 - \tau) = v'(y_t/n_t).$$

Given the parameterization $v(x) = \frac{\eta}{\eta+1} x^{\frac{\eta+1}{\eta}}$, this may be rewritten as

$$y_t = \left(\frac{\lambda}{1 - \delta(1 - \lambda)(1 + g)} \cdot (1 - \tau) \right)^{\eta} n_t.$$

This is equation (14) in Proposition 5. Finally, note that the solution for (y_t, \hat{z}_t) characterized above satisfies the property $y_s = y_t(1 + g)^{s-t}$ on which the derivations relied.⁴⁰

B.6 Endogenous λ

If effort is observable or if workers can commit to an effort level, equilibrium earnings equal $y_t = (1 - \tau)^{\eta} n_t$ as in a standard model. This maximizes worker-firm surplus (efficiency). We consider instead a setting where effort is unobservable without costly performance evaluations of workers. Evaluating a given worker costs q and reveals true effort y_t in the current period. Evaluations are carried out randomly with frequency λ . Considering a steady state with constant productivity n and effort y (to simplify exposition), we solve for the constrained-efficient solution of (y, λ) that maximizes worker-firm surplus.⁴¹ The per-period surplus is given by

$$S = (1 - \tau) [y - q\lambda] - nv(y/n),$$

where the term in square brackets is the net output/income generated. Note that, in this specification, we assume that evaluation costs $q\lambda$ are tax deductible. This will be the case if, for example,

⁴⁰It is possible to allow for stochastic productivity shocks. For example, we can show that Proposition 5 still applies if we assume $n_t = \tilde{n}_t \mu_t$, where $\tilde{n}_{t+1} = \tilde{n}_t(1 + g)$, $\mu_{t+1} = \mu_t + \epsilon$, and ϵ is white noise.

⁴¹The solution can be decentralized in a competitive economy where workers receive compensation $(1 - \tau)(y - f)$ where f equals $q\lambda$, which corresponds to firm spending on worker evaluations. In this situation, firm profits are zero in equilibrium.

the costs of performance evaluations reflect labor costs.

The solution to y is still given by (4). The first-order condition for λ equals

$$\frac{dS}{d\lambda} = [1 - \tau - v'(y/n)] \frac{dy}{d\lambda} - q(1 - \tau) = 0.$$

With costless verification ($q = 0$), we have $v'(y/n) = 1 - \tau$. Given the parameterization $v(x) = \frac{\eta}{\eta+1} x^{\frac{\eta+1}{\eta}}$ used previously, this implies $y = (1 - \tau)^\eta n$ and is implemented by setting $\lambda = 1$ according to equation (4). With costly verification ($q > 0$), the incomplete information creates a wedge between the marginal benefit of effort $1 - \tau$ and the marginal cost of effort $v'(y/n)$.

By inserting the marginal disutility of effort and using equation (4), we may rewrite the optimality condition as

$$\frac{dS}{d\lambda} = (1 - \tau) \frac{(1 - \lambda)(1 - \delta)}{1 - (1 - \lambda)\delta} \cdot \frac{dy}{d\lambda} - q(1 - \tau) = 0.$$

By differentiating equation (4) and rearranging terms, we obtain

$$\frac{dy}{d\lambda} = \frac{\eta(1 - \delta)}{\lambda(1 - (1 - \lambda)\delta)} y,$$

which may be inserted into $dS/\partial\lambda = 0$ to arrive at the following equilibrium condition for λ :

$$\frac{\lambda}{1 - \lambda} = \frac{\eta(1 - \delta)^2}{\gamma(1 - (1 - \lambda)\delta)^2}, \quad (30)$$

where $\gamma \equiv q/y$ denotes the evaluation cost in proportion to output. We may interpret γ as capturing the degree/cost of imperfect information, which determines where λ lies in the interval between perfect verification ($\lambda = 1$ which obtains when $\gamma = 0$) and no verification ($\lambda = 0$ which obtains when $\gamma = \infty$). In general, for a positive and finite value of γ , the evaluation frequency λ lies between 0 and 1, thereby giving rise to the dynamic return mechanisms characterized in this paper. As for comparative statics, equation (30) shows that λ is decreasing in the evaluation cost γ , increasing in the effort elasticity η , decreasing in the discount factor δ , and independent of τ . The last result relies on the (natural) assumption that evaluation costs are tax deductible.