

Discrete Earnings Responses: Empirical Evidence and Implications

Tuomas Kosonen and Tuomas Matikka*

December 31, 2019

Abstract

We provide new evidence of discrete earnings responses to tax incentives, and study the broader implications of discrete rather than continuous earnings adjustment. We utilize an income notch created by the study subsidy system for university students in Finland and a reform that shifted out the location of the notch to uncover the mechanisms behind earnings adjustments. We find clear evidence of discrete earnings responses, revealing that wage earners even in the part-time labor market can face significant restrictions in their available earnings choices. Our simulation results highlight that the canonical continuous-choice labor supply framework can provide biased earnings elasticity estimates when the actual earnings adjustment process is discrete.

Keywords: discrete earnings choices; labor supply; tax elasticity

JEL Classification Codes: H21, H24, J22

Acknowledgments: We thank Raj Chetty, Alexander Gelber, Jarkko Harju, Pat Kline, Wojciech Kopczuk, Jim Poterba, Emmanuel Saez, Danny Yagan and Gabriel Zucman, as well as participants in numerous conferences and seminars for their helpful comments and discussions. We are also grateful to Spencer Bastani for helping us getting started with the simulation model. Funding from the Academy of Finland (No. 303689) is gratefully acknowledged.

*Kosonen: Labour Institute for Economic Research (Helsinki, Finland), Academy of Finland and CESifo, tuomas.kosonen@labour.fi. Matikka: VATT Institute for Economic Research (Helsinki, Finland), Academy of Finland and CESifo, tuomas.matikka@vatt.fi.

1 Introduction

In all developed countries, a major part of the population receive their income through wage earnings that are subject to income taxes. Therefore, a question of how taxpayers' earnings respond to taxes has received extensive attention in the economics literature (see Saez *et al.* (2012) for a survey). Recently, many studies have analyzed factors that might prevent individuals from responding to taxes, commonly labeled as optimization frictions (see e.g. Chetty *et al.* (2011), Chetty (2012), and Kleven and Waseem (2013)). These are typically modeled as small deviations from the canonical continuous-choice labor supply model that hinder individual responses to taxes. A more fundamental deviation from the standard framework is to assume that individual earnings choices comprise of a discrete set of available choices instead of continuous earnings adjustment, a feature that some previous studies have analyzed theoretically and in structural labor supply applications (see e.g. Dickens and Lundberg (1993), Saez (2002), and Kreiner *et al.* (2015)). Discreteness in earnings choices could arise from various sources. For example, working hours restrictions and other labor market regulation such as minimum wages are typical in many countries, and employment contracts with specific job descriptions and long notice periods further limit continuous earnings adjustments to tax policy. Also, workers with specific training and occupational skills rarely face a large set of available jobs that match their skills and preferences (Saez, 1999).

Despite the fact that discrete earnings choices are likely to pose a relevant constraint for a significant proportion of wage earners, they have received only limited attention in the literature. Specifically, quasi-experimental evidence on prevalence of discrete earnings responses is scarce. Moreover, the consequences of a discrete rather than continuous earnings model on reduced-form tax elasticity estimation are not thoroughly explored. In this paper, we fill these important gaps in the literature by providing novel and transparent evidence of discrete earnings supply among wage earners. We show that the standard differences-in-differences (Gruber and Saez, 2002) or bunching method (Saez, 2010) tax elasticity estimates relying on the canonical model that assumes continuous earnings supply can be significantly downwards biased if the actual earnings adjustment process is discrete. As the elasticity of taxable income (ETI) is a key measure for welfare losses (Chetty, 2009), downward-biased elasticity estimates could lead to an assessment that income taxation induces smaller welfare losses than it actually does.

We utilize a novel empirical design to study the hypothesis of discrete earnings responses: a combination of an income notch creating strong local tax incentives, and a reform that shifted out the location of the notch. The institutional setting we utilize is a monthly study subsidy for Finnish higher education students. A university student loses eligibility for one month of the subsidy (approximately 500 euros) if her earnings exceed a predetermined annual gross income threshold (9,200 euros before 2008). Therefore, the

income threshold creates an income notch, a strong local tax incentive for students above which their disposable income reduces sharply. As Finnish university students typically participate in flexible part-time and temporary labor markets during their studies both within and between semesters, this notch creates a binding budget constraint for a majority of students. In 2008, the location of the threshold was increased by 30%, allowing students to earn more income before they lose the subsidy they are eligible for. Using the relocation of the notch, we can detect responses to a local incentive from a broader income range, enabling us to uncover whether the notch affects the earnings choices far away from the discontinuity. In the absence of the reform, we could not identify those who are affected by the notch even when they are located further away from it. In addition, the extensive register-based panel data covering all Finnish taxpayers enable us to follow the same students over time and link their earnings to the characteristics of the firm they worked for, allowing us to shed more light on the mechanisms at play.

As our main empirical result, we find that the change in the location of the notch caused distinctive earnings responses in a broad income range. We find that many students who were located far below the old notch before the reform significantly increased their earnings precisely at the time when the location of the threshold was increased. As a result, the overall shape of the earnings distribution changed within the reform, as the density at lower incomes reduced, and consequently, increased at higher income levels. To capture the full earnings response to the relocation of the notch, we develop a novel method that calculates the change in the relative density in the distribution associated with the reform. The method operates in the same spirit as the bunching method, which estimates the local excess mass associated with a kink or a notch, but in this case the estimation covers the earnings distribution from a wider range.

We account for potential changes in earnings that occurred for other reasons than the reform by utilizing a differences-in-differences approach and a control group consisting of non-student young wage earners. These workers are in the same part-time labor markets and have similar earnings distribution as the student population but they did not face any changes in their incentives. We do not detect any changes in the shape of the earnings distribution of our control group. This indicates that the earnings responses of students did not arise from other contemporary changes in the part-time labor market in Finland. Also, we find that students appeared to be aware of the reform. A visible local excess mass immediately below the old notch completely disappeared, and a new but somewhat smaller excess mass appeared just below the new location of the notch right after reform.

In addition, we provide panel data evidence showing that large and discrete individual-level earnings increases are much more common at the time of the reform compared to the period before it. For example, in 2007–2008, students originally below the income threshold were much more likely to increase their earnings by 50% or more, compared to the period before the reform. These results further support the notion that the earn-

ings responses within a broad income range arise because of individual-level relocation responses rather than, for example, extensive margin responses stemming from more productive students participating in the labor market after the reform. Therefore, our results show that even wage earners in temporary and part-time labor markets appear to face significant earnings constraints. Any frictions affecting the earnings responses of these types of workers are also likely to be present for other wage earners in more permanent and arguably more discrete labor markets (Saez, 1999).

The next important step is to consider in more detail which mechanisms the empirical results outlined above suggest are at play here. We utilize simple theoretical arguments to discuss that the most likely explanation for the earnings responses within a broad income range well below the notch is that the earnings choices are constrained by a set of discrete earnings locations. First, the canonical labor supply model with continuous earnings adjustment cannot explain the observed pattern. In that model, individuals who are originally located far below the notch are unaffected by the reform, as they are assumed to have already optimized their earnings location as a response to the existing tax system which does not change at these income levels when the location of the notch is increased.

Instead, in the discrete earnings choice model, an individual considers multiple earnings locations in the income distribution, justifying large earnings responses to the reform from below the notch. For example, one available earnings location for an individual could be far below the original notch and the next one above it. When the location of the notch is increased, an individual compares her utility in the two locations. If the location at a higher income level becomes more attractive after the relocation of the notch, an individual originally far below the old notch responds to the reform with a large discrete jump in earnings. More generally, with heterogeneous earnings locations across individuals, a certain fraction of individuals far below the original notch respond to the reform by increasing their earnings, leading to an overall shift in the earnings distribution. This is exactly what we observe in our empirical analysis.

Moreover, we provide illustrative simulation results showing that the continuous-choice model, even when augmented with adjustment frictions and optimization errors discussed in the recent literature (see Chetty (2012) and Kleven and Waseem (2013)), does not produce earnings responses from below the notch when the threshold is relocated. In contrast, when the available earnings choices are limited to a discrete set in the model, the resulting earnings responses and the shape of the distributions resemble our empirical results. Thus, our empirical results support the discrete earnings model analyzed theoretically and in structural applications in the previous literature (Dickens and Lundberg 1993; van Soest 1995; Saez 2002; Kreiner *et al.* 2015; Beffy *et al.* 2019).

We find further reduced-form support for the discrete earnings hypothesis by examining the responses of two specific subgroups of students: those who work in arguably more

discrete labor markets (public sector, or research, manufacturing and construction in the private sector) and those working in less discrete labor markets (restaurants, bars and cafes, hotels, cleaning and security services). The latter group has typically more flexible working hours and are subject to hourly rather than monthly wages, compared to the first group. We find a significantly larger shift in the earnings distribution from below the notch for students who work in more inflexible and discrete labor markets, compared to students working in less discrete labor markets. This finding supports our hypothesis that earnings responses from a wider income range are caused by discrete earnings choices.

We also estimate the welfare losses created by taxes in our empirical setting and income taxation in general. First it is good to note that when earnings choices are discrete, a correct elasticity of taxable income concept is the mobility elasticity. That accounts for the fraction of people responding to changes in average tax rates between two distinct locations instead of the change in income in response to change in marginal tax rate as the conventional elasticity of taxable income does (Saez 2002; Kreiner *et al.* 2015). Because in the discrete model individual who are not directly affected by tax changes can still respond to them, we need to apply a tax change to these individuals to calculate the total response to a tax reform. The mobility elasticity achieves this by calculating the change in average tax rate between the location individual is before the tax change and where the individual would want to be located after the change. This is the key difference between the normal elasticity of taxable income concept that applies a change in marginal tax rate only to those who are directly affected by the tax reform based on their current location in the tax schedule.

We apply a novel method to estimate the mobility elasticity for students using the relocation of the notch. In the estimation, we relate the average earnings response to the changes in net of average tax rates across the average earnings locations the students are located before and after the reform. We estimate a mobility elasticity of 0.18, which is close to earlier earnings elasticity estimates in the literature (see Saez *et al.* (2012) for a survey) and also close to what Sogaard (2019) finds for university students in Denmark for a similar setting involving a kink instead of a notch.

Moreover, discrete earning choices can entail significant broader implications for reduced-form tax elasticity estimation. Most of the existing empirical literature on the elasticity of taxable income assume some version of a standard continuous-choice model, either explicitly or implicitly (Saez *et al.* 2012; Kreiner *et al.* 2015). We use our simulation model to analyze how the observed elasticity estimate depends on the underlying modeling assumptions, focusing on the implications when the underlying earnings adjustment process is discrete rather than continuous. We simulate a simple tax reform where an increase in the marginal tax rate from 40% to 60% above a predetermined kink point is repealed, representing a tax rate cut for higher incomes. We then apply a differences-in-differences approach to estimate tax elasticity using the standard assumptions in the

literature and the mobility elasticity estimator which correctly estimates the tax elasticity when earnings choices are discrete.

When comparing these two approaches, we find that the standard estimation procedure produces a significant negative bias to the estimate whenever individual earnings choices are sufficiently discrete. This bias comes from the fact that the standard approach would typically use individuals who do not face a tax rate change based on their pre-reform earnings as a control group for those who do, as in a seminal paper by Gruber and Saez (2002). In our example, this would mean that individuals originally below the kink point would form a control group for those who were above the kink and faced a tax rate cut based on their pre-reform earnings. However, when earnings choices are discrete, also individuals below the original kink can respond to this reform if their next available earnings location is above the kink point. Therefore, this control group is affected by the reform, causing the estimates derived using the standard approach to be downward biased.

Thus, these findings imply that the interpretation of the ETI estimates in earlier studies can be very sensitive to the underlying labor market constraints and the earnings adjustment mechanisms. Nevertheless, this type of a bias is not relevant for all reduced-form elasticity estimates, such as in the literature studying the impact of the earned income tax credit (EITC) on earnings where women without children are typically used as a control group for women with children who are eligible for larger tax credits (see Kleven (2019) for a recent review). In these types of set-ups, the control group is typically unaffected by the changes in taxes occurring for the treatment group, implying that potential discrete earnings responses do not bias the estimator.

Furthermore, irrespective of the estimator used, the discrete choice constraint typically leads to smaller earnings elasticity estimates compared to the continuous earnings scenario. This implies that discrete earnings choices provide a feasible explanation for why we typically observe rather modest ETI estimates in the literature (see e.g. Saez *et al.* (2012) and Neisser (2019)). Nevertheless, with some parameter values the average elasticity can even be higher in the discrete model than in the continuous model, particularly when earnings choices are less discrete and the original kink point is located at a lower earnings level in the distribution.

The magnitude of the tax elasticity relates to the discussion on the welfare implications of taxes (see e.g. Chetty (2009, 2012)). Under the assumption that discrete earnings choices create a persistent constraint for individuals, the observed mobility elasticity estimate represents the sufficient statistic for welfare analysis. In contrast, under more temporary adjustment frictions discussed in the earlier literature, such as inattention, low tax salience and optimization errors, the empirically observed elasticity and the structural elasticity would differ such that the observed empirical estimate would be smaller (Chetty *et al.* 2009; Chetty *et al.* 2011; Chetty 2012; Chetty *et al.* 2013; Chetty and Saez 2013;

Kleven and Waseem 2013; Gelber *et al.* 2018; Gelber *et al.* 2019; Søggaard (2019)). Thus, the distinction between our paper and the earlier studies is that these types of frictions have clearly different welfare implications compared to discrete earnings choices. This paper and the earlier literature show together that both adjustment frictions and discrete earnings choices are likely to provide a relevant explanation for modest observed responses to taxes in the literature.

Furthermore, we contribute to the bunching literature that estimates responses to local discontinuities in various types of incentives, summarized by Kleven (2016). In particular, we add to the recent literature discussing the limitations of the bunching method in measuring the elasticity of taxable income (see e.g. Blomquist *et al.* (2018)). We find that when earnings choices are discrete, the local bunching method significantly underestimates the true response to the notch, due to the fact the local approach cannot capture earnings responses from a broader income range below the notch. Consequently, in our empirical example, the bunching method underestimates the associated earnings elasticity by a factor of 2.5 when compared to the mobility elasticity estimate.

Finally, we contribute to the literature discussing discrete earnings responses. Most of this literature is theoretical, including, for example, Dickens and Lundberg (1993) and Saez (1999, 2002). In addition, the structural labor supply literature often assumes that working hours decisions or job changes are discrete in nature (see van Soest (1995), Kreiner *et al.* (2015), Beffy *et al.* (2019), and Löffler *et al.* (2018) for a recent review). However, there is only limited quasi-experimental evidence of the prevalence of discrete earnings responses, perhaps due to a lack of suitable quasi-experimental variation and data for identifying the mechanisms behind the earnings adjustment process. One exception is Blundell *et al.* (2008), who estimate the intensive-margin labor supply responses of single mothers to changes in various in-work benefit programs in the UK. They find that the responses are governed by discrete working hours responses between jobs rather than continuous labor supply or wage rate adjustments. We contribute to this literature by providing novel and transparent evidence of significant discrete earnings responses among wage earners, combined with illustrative simulation results highlighting the key role of discrete earnings choices in explaining our findings.

This paper proceeds as follows: Section 2 presents the relevant institutions and empirical methods. Section 3 presents the main results. In Section 4, we discuss the theoretical mechanisms, and Section 5 presents our simulation results and discusses the broader implications of discrete earnings choices. Section 6 concludes.

2 Institutions, data and empirical methods

2.1 Study subsidy for university students

In Finland, all students who are enrolled in a university or polytechnic can apply for a monthly-based study subsidy, administered nationally by the Social Insurance Institution of Finland (hereafter SII). The subsidy is intended to enhance equal opportunities to acquire higher education, and to provide income support for students who often have low disposable incomes. In Finland, university education is publicly provided and there are no tuition fees. A large proportion of individuals receive higher education in Finland (approximately 40% of individuals aged 25-34 have a degree), and the study subsidy program is widely used among students.

The maximum amount of the subsidy was 461 euros per month in 2007. The default number of subsidy months per year is 9, which the bulk of the students also choose. The eligibility for the study subsidy depends on personal annual gross income (labor income + capital income), and completing a certain predefined number of credit points per academic year. Parental income or wealth does not affect eligibility nor the amount of the benefit for students not living with their parents.¹

The discontinuity in labor supply incentives is created by an income threshold. If annual gross income is higher than a predetermined threshold, the study subsidy of one month is reclaimed by the SII. This results in an increase in the average tax rate, or an increase in the implied marginal tax rate of over 100%, in a region just above the threshold, creating a *notch* in the budget set of students. With the default 9 months of the subsidy, the annual income threshold was 9,260 euros in 2007. An additional month of the subsidy was reclaimed for an additional 1,010 euros of income above the threshold. This implies that the schedule ultimately comprises of multiple notches. Students can deviate from the default of 9 months and alter the number of subsidy months by application, or by returning already granted subsidies by the end of March in the next calendar year. Having more study subsidy months reduces the income threshold, and vice versa.²

The study subsidy program was reformed in 2008. The main outcome of the reform was that the income threshold was increased by approximately 30%. For a typical student with 9 study subsidy months, the annual income threshold increased from 9,260 to 12,070 euros. In addition, the monthly study subsidy was slightly increased from 461 to 500 euros per month. As with the old regime, an additional month of the subsidy is reclaimed after

¹The full study subsidy includes a study grant and a housing benefit. The standard study grant was 259€/month and the maximum housing benefit 202€/month in the academic year 2006/2007. Housing benefits are granted only for rental apartments, and the housing allowance is reduced if spousal gross income exceeded 15,200 per year (in 2007). In addition to the study subsidy, students can apply for repayable student loans which are secured by the central government.

²In 2007, the formula for the annual income threshold was the following: 505 euros per study subsidy month and 1,515 euros per month without the study subsidy, plus a fixed amount of 170 euros.

an additional 1,310 euros of gross income above the threshold.³ Other details of the system were not changed, including the academic criteria.

Figure 1 illustrates the study subsidy schedule before and after 2008 for a student who collects the default 9 subsidy months. First, the figure shows that students face large local incentives not to exceed the income threshold. Once the income threshold is exceeded, losing one month of the subsidy causes a significant dip in disposable income. Therefore, the study subsidy notch induces a strictly dominated region just above the threshold where students can earn more disposable income by reducing their gross earnings. Furthermore, the figure underlines the distinctive change in incentives caused by the increase in the income threshold in 2008, highlighting that the reform encouraged to increase earnings above the old income threshold. Finally, Table A1 in Appendix A shows the income thresholds in numbers before and after 2008, and presents the relative loss in disposable income that incurred if the income threshold was exceeded.

2.2 Data and descriptive statistics

Although the majority of students have access to the study subsidy and repayable student loans, most university students in Finland also work part-time during studies within and between semesters. Therefore, the means-testing of the study subsidy creates a binding budget constraint for a majority of students. In our analysis, we use panel data on all working-age individuals (15–70 years) living in Finland in 1999–2013, provided by Statistics Finland. These data include a variety of register-based variables, such as detailed information on tax register and social benefit items, including information on the study subsidy program. With these data, we can analyze responses to the incentives created by the program and learn how various individual characteristics affect behavioral responses.

Table 1 shows the descriptive statistics for all students in 1999–2013. Average annual labor income among all students is 9,130 euros. We observe that 77% of students earned more than 500 euros of labor income in a year. In addition, less than 60% of students received labor earnings from only one employer, suggesting that students tend to work in different types of jobs during the year.

Overall, these observations indicate that many students work in part-time or temporary jobs during their studies and breaks between semesters in order to increase their disposable income and/or to gain work experience while studying. Also, 18% of students work in manufacturing (including construction), 15% in the service sector (mainly restaurants and hotels), and 37% in administrative and support services or in the public sector.

³After 2008, the formula for the annual gross income limit was the following: 660 euros per study subsidy month and 1,970 euros per month when no study subsidies are collected, plus a fixed amount of 220 euros.

The average number of study subsidy months collected per year is 6.7. The share of students receiving the default subsidy of 9 months is 32%. The group receiving the default subsidy is very similar to the overall student population, expect that they are, on average, slightly younger (22.4) and have less labor income (5,633 euros) than all students. An average student in the data has been studying for approximately 2 years. Finally, 13%, 16% and 30% of students in our data study arts and humanities, business and social sciences, and technology or health and social services, respectively.

In the forthcoming analysis, we focus on students who received 9 months of study subsidy before and after 2008. For this group, the income threshold increased from 9,260 to 12,020 euros. This restriction is not very selective as a bulk of students receive 9 months of the study subsidy, partly because it is the default choice and partly because it presumably creates a good balance between subsidies and labor earnings for many students. The advantage we gain by fixing the number of subsidy months is that we can isolate the effect of the change in the location of the threshold on the earnings distribution for a large part of the student population. In addition, we restrict our sample to students who do not graduate within a given year in order to avoid the effects of potential earnings shocks after graduation. However, dropping graduates does not affect the main results in any significant way.

2.3 Estimation

The income threshold reform creates a unique empirical set-up to disentangle different types of earnings responses to a large and salient change in tax incentives. In our analysis, we are particularly interested in investigating whether local tax incentives, such as notches, affect income distributions further away from the local discontinuity in incentives. Thus, we examine the shape and location of the whole income distribution before and after the reform, and develop a new method to estimate these changes building on the local bunching method.

Behavioral responses to local discontinuous changes in the budget set, such as tax rate kinks or notches, are in the recent literature predominantly estimated using the bunching method (see Kleven (2016) for a summary). In the local bunching approach, the behavioral response caused by the threshold is estimated by relating the observed excess mass in the earnings distribution just below the threshold to the counterfactual density that would exist in the absence of the discontinuity. Typically, the counterfactual density is estimated by fitting a flexible polynomial function to the observed earnings (z) distribution, excluding an income range $[z_L, z_H]$ of the observed distribution just around the threshold (z^*). Graph (a) in Figure 2 illustrates local bunching below a tax notch in a hypothetical earnings distribution. We discuss the local bunching approach in further detail in Appendix B.

However, the local bunching method could produce biased estimates of the extent of behavioral responses if the threshold affects the earnings distribution further away from the threshold. In this case, the local bunching estimate does not sufficiently capture the distortions caused by the discontinuous change in incentives. One potential cause for responses within larger income intervals is constraints that limit the possibility to continuously adjust earnings. Under this constraint, individuals can adjust their earnings only in a discrete manner, in contrast to continuous earnings adjustments assumed in the bunching model (Saez 2010; Kleven and Waseem 2013).

In order to detect and estimate discrete earnings responses, we follow the baseline idea of the local bunching method but evaluate the effects of the study subsidy income threshold on the overall shape of the earnings distribution further away from the notch. In the analysis, we exploit the 2008 threshold reform and the pre-2008 earnings distribution as a counterfactual when numerically characterizing changes in the distribution caused by the increase in the income threshold after 2008. We denote the distributions in relative terms in order to take into account the fact that the number of students at certain income levels might slightly differ between the years.⁴

Graph (b) in Figure 2 illustrates the estimation of broader changes in the earnings distribution of students. We set the lower limit z_L well below the threshold in order to capture the broader changes in the distribution, in contrast to local bunching method that focuses on responses just below the notch. More formally, the change in the shape of the overall distribution below the location of the old threshold can be characterized as

$$\hat{b}(z) = \frac{\sum_{i=z_L}^{z_H} [(c_j^B/N^B) - (c_j^A/N^A)]}{\sum_{i=z_L}^{z_H} (c_j^A/N^A)/N_j} \quad (1)$$

where $\sum_{i=z_L}^{z_H} (c_j^k/N^k)$ is the relative share of students within an income range $[z_L, z_H]$, and $k = B, A$, where B denotes the time period before the reform and A after 2008. c_j is the count of individuals in income bin j , and z denotes the income level in bin j , and N^k denotes the overall number of students and N_j denotes the number of bins within $[z_L, z_H]$.⁵

Furthermore, to estimate the causal impact of the reform on changes in the earnings distribution, we develop a novel differences-in-differences estimator. In general, other factors than the change in the location of the notch might inflict changes in the shape of the earnings distribution, such as overall changes in the economic environment and the labor market. In order to take these issues into account, we utilize the changes in the

⁴In the standard cross-sectional bunching analysis, using relative distributions instead of frequency distributions produces identical estimates of the relative excess bunching at the discontinuity.

⁵Our approach is similar in nature to other methods used to analyze changes in density distributions (ADD CITATION HERE). However, by extending the methods used in the local bunching approach, we can more feasibly relate our estimates to those derived using the local bunching method that has been predominantly used to analyze behavioral responses to local discontinuities in incentives such as income notches (see Kleven (2016)).

distribution of young, part-time non-student workers who match students' job and age characteristics. Even though current students might differ from current non-student part-time workers in some relevant non-observed characteristics, the income development of other part-time workers still captures the underlying general economic trend that affects the overall earnings potential of the part-time labor force, which we thus aim to control for using the differences-in-differences approach.

The non-student part-time workers included in the analysis are not subject to the income threshold, but are of the same age as students and work in similar types of jobs. This group thus resembles the treated students as they have similar average labor earnings as students and many of these workers have more than one employer within a year, similarly as the student population. The characteristics of young, part-time workers are described in Table 2.⁶ However, our differences-in-differences estimator does not require that the two groups have exactly the same pre-reform income distribution or characteristics. As in the standard differences-in-differences method, the identification assumption is that the two distributions should develop similarly over time in the absence of any changes in incentives.

In this estimation, we follow the approach described above to calculate the change in the density of the earnings distribution between two time periods, and subtract the change in the non-student part-time workers' earnings distribution from the change in the students' distribution

$$\hat{b}_d(z) = \left[\frac{\sum_{i=z_L}^{z_H} [(c_j^B/N^B) - (c_j^A/N^A)]}{\sum_{i=z_L}^{z_H} (c_j^A/N^A)/N_j} \right]^S - \left[\frac{\sum_{i=z_L}^{z_H} [(c_j^B/N^B) - (c_j^A/N^A)]}{\sum_{i=z_L}^{z_H} (c_j^A/N^A)/N_j} \right]^P \quad (2)$$

where superscript S denotes students and P non-student part-time workers. This estimate thus summarizes the broader change in the earnings distribution of students caused by the reform while taking into account other potential changes in the labor market environment of part-time workers.⁷

⁶The group of non-student part-time workers is selected to roughly match students' job and age characteristics. Students typically work in part-time jobs or in full-time jobs for a part of the year, i.e. they work less than 12 months a year. In addition, students tend to be young. Thus, the control group comprise of individuals who we observe to have less than 12 working months per year, and who are 19–24 years old. The age interval is chosen to match between the 25–75 percentile points of the students age distribution. Our results are not sensitive to small changes in the composition of the non-student group.

⁷Following the bunching literature, the standard errors for $\hat{b}_d(z)$ are calculated using a residual-based bootstrap procedure (see e.g. Kleven and Waseem (2013)). First, we fit a flexible polynomial function to both pre- and post-reform relative earnings distributions of students and other young part-time workers. We then generate a large number of new estimates for the distributions by randomly re-sampling the residuals from these regressions (with replacement). The standard error is defined as the standard deviation of $\hat{b}_d(z)$ based on the bootstrapped distributions.

3 Main results

Figure 3 shows the labor earnings distributions of students and non-student part-time workers within an income range of 0–18,000 euros in 2006–2007 and 2008–2009, denoting the pre- and post-reform years, respectively.⁸ Remarkably, the figure shows that the earnings distribution of students has a significantly different shape in 2008–2009 than before the reform, as the earnings have increased in a wide income range also below the old location of the income threshold. After the reform, the income distribution shifted to the right from a large region below the threshold, from about earnings of 2,000 euros onward. Contrary to students, the earnings distribution of non-student part-time workers remained practically constant between 2006–2007 and 2008–2009, indicating no other contemporary changes in earnings among other young part-time workers who are not subject to the income threshold nor changes in its location.

To quantify the changes in the distribution, we estimate the differences-in-differences equation (2) within an income range of 0–9,200 euros, thus including the whole distribution below the old income threshold. The estimate is large (9.809, with a standard error 1.01), suggesting that the magnitude of the change in the overall earnings distribution is economically and statistically significant. This estimate is over three times larger than the standard local bunching estimate, 2.931 (0.875), estimated following the reduced-form method of Kleven and Waseem (2013) within an income range just below the threshold (8,100–9,200 euros) before the reform.⁹ In order to further characterize the general magnitude of the overall income response, we estimate an average earnings increase of 550 euros per student when accounting for the overall changes in the shape of the earnings distribution, which corresponds to a roughly 10% average increase in labor earnings.

In addition, Figure 3 shows that at least some fraction of students are aware of the location of the income thresholds and are able to precisely adjust their labor earnings to them, as local bunching just below the threshold is significant and clearly visible both before and after 2008. Furthermore, the local bunching response disappeared below the old threshold immediately after the reform, and a new excess mass appeared just below the new threshold right after the reform. Therefore, we find no evidence of some students still believing that there is a notch at the old location nor that there would be a sluggish local response to the relocation of the threshold. Furthermore, even though the study subsidy schedule ultimately consists of multiple notches, we observe a distinctive local response only to the first income threshold.¹⁰

⁸The figure includes only labor earnings as receiving capital income is very rare among university students and other part-time workers.

⁹The local bunching method and local bunching results are discussed in more detail in Appendix B.

¹⁰Additional examination of excess bunching before and after the reform reveals, as further illustrated in Figure B2 in Appendix B, that bunching is slightly larger before the reform than after it. One intuitive explanation for this finding is that local incentives not to exceed the notch are somewhat smaller after 2008, since the relative significance of losing one month's subsidy in terms of disposable income is now

As a first robustness check, Figure 4 plots students' earnings distributions from a longer time period before and after 2008. The figure shows that the change in earnings occurred exactly at the time of the relocation of the income threshold, indicating that any gradual shifting of the earnings distribution does not explain the observed pattern. In addition, Figure A1 in Appendix A shows the distributions in 2006–2007 and 2008–2009 when we re-weight the student population in the latter period to match the pre-reform characteristics of students in terms of age, field of study and field of industry. This bin-level inverse probability weighting procedure accounts for potential changes in key student characteristics over time. However, re-weighting does not change the outcomes in a significant manner, indicating that potential changes in the characteristics of the student population are not likely to explain the results either.¹¹

Next, we present more detailed evidence of the discrete earnings responses of students utilizing the panel dimension of our data. Overall, these results show that many students responded to the relocation of the income threshold with a large increase in their income instead of marginal earnings adjustments across the distribution. First, graph (a) of Figure 5 presents the likelihood of increasing individual earnings by 50% or more relative to base-year income. We observe that large increases in earnings are significantly more likely when the threshold was increased compared to previous years. For example, the prevalence of annual earnings increases larger than 50% doubled from 5% to 10% in income bins below the old income threshold at the time of the reform. In contrast, there are no significant differences in large earnings increases between the pre- and post-reform years in bins above the new income threshold.

Second, graph (b) of Figure 5 shows that the likelihood of locating above the old income threshold in the next year increased significantly in the bins below the new threshold at the time of the reform, compared to the years prior to 2008. Again, the fact that the likelihood of being located above the old notch increased in income bins far below the old threshold illustrates that a notable share of students responded to the reform with a large increase in their earnings when their budget constraint was relaxed at higher earnings levels.

Third, in graph (c) of Figure 5, we analyze individual-level earnings responses in further detail. The figure presents the average individual-level changes in real labor

smaller than before 2008 when the threshold was at a lower income level.

¹¹In addition, we have studied other potential changes at the time of the reform that might affect observed changes in the shape of the earnings distribution. First, there were no significant changes in the distribution of subsidy months associated with the reform, and 9 months is the most typical choice in all of the years around the reform. This indicates that current students responded to the reform by changing their earnings, but not, on average, by claiming more or less subsidies per year. Second, we looked at whether the reform is accompanied by extensive margin responses, but the share of students not working at all (earning less than 500 euros per year) did not change significantly at the time of the reform. Therefore, these types of responses do not explain the change in the shape of the observed earnings distributions around the 2008 reform. These results are not reported but available from the authors upon request.

income in 2005–2006, 2006–2007 and 2007–2008. Overall, the figure shows that average changes in individual income are very similar in the years before the reform, and that there is a visible pattern of mean reversion (on average, starting from a low income level leads to larger income in the next year, and vice versa). The figure shows that labor income increased significantly in 2007–2008 compared to the years before the reform for students below the new income threshold. This pattern is observable even for students with base-year earnings around 3,000–6,000 euros, which is well below the old threshold. However, we find no significant difference between the years for income bins above the new threshold, suggesting that the rapid increase in earnings below the old threshold stems from the change in the location of the income threshold.

Overall, we find clear evidence that the 2008 threshold reform induced large earnings responses for students who were previously located well below the old income threshold, consistent with the hypothesis that the relaxed budget set constraint created large and discrete earnings responses for many students. This observation is particularly surprising as we are studying the student population, who typically work in flexible part-time or temporary jobs, and are thus likely to have a variety of opportunities to adjust their labor supply and earnings. In other words, any frictions related earnings adjustment mechanisms stemming from the labor market are presumably much less relevant for this population compared to regular wage earners.

4 Conceptual framework and implied mechanisms

4.1 Earnings choice models

Next, we discuss in more detail which theoretical mechanisms the empirical results outlined above suggest are at play here. First, we briefly outline the general framework for analyzing the welfare effects of taxes and income transfers. A key method in measuring the welfare consequences of various policies is the sufficient statistics approach (see e.g. Chetty (2009)), where the idea is to estimate a well-identified reduced-form parameter that provides a direct measure for the welfare loss. A prime example of the sufficient statistics approach is the extensive elasticity of taxable income (ETI) literature, where the ETI with respect to the marginal tax rate directly delivers the sufficient statistic for the welfare analysis of income taxes (see Feldstein (1999) and Saez *et al.* (2012) for a review). Following Chetty (2009) and Slemrod (1998), the welfare (W) effect of a tax rate change can be formalized as follows

$$\frac{dW}{d\tau} = \frac{dR}{d\tau} - \frac{\partial R}{\partial \tau} \Big|_M \quad (3)$$

where $dR/d\tau$ denotes how total tax revenue (aggregate earnings over all individuals times the tax rate) responds to a change in the tax rate, and $\partial R/\partial\tau|_M$ denotes the mechanical change in revenue absent any behavioral responses. Therefore, the welfare loss stems from how much (taxable) earnings respond to a change in the tax rate, i.e. the ETI.

As discussed in Chetty (2009), the underlying earnings response mechanism does not change the basic idea of the sufficient statistic formula. For example, assuming discrete rather than continuous earnings choices does not change the term referring to the impact on government revenue, as it is defined similarly in both cases. However, the elasticity concept defining the extent of the behavioral response and the scope of the welfare effect depends on the chosen model and estimation method.

The aim of the remainder of this section is twofold: first, we discuss which earnings choice models are or are not compatible with our empirical results presented above. The main feature we want to explain is the shifting of the income distribution from a wide income range below the old threshold following the change in the location of the notch. Then, we describe how the different modeling choices discussed in the literature affect the capability of the model to capture a consistent earnings elasticity compatible with large earnings responses from a broad income range.

We start with a canonical continuous earnings supply model and then extend that with adjustment frictions and discrete earnings choice sets. A continuous-choice framework features a utility function over consumption and leisure and a linearized budget set consisting of earnings, consumption and income taxes. The utility function is $u(c, z)$, where c denotes consumption (net earnings) and z gross earnings, and $u_c > 0$ and $u_z < 0$. The budget set is described as $c = (1 - \tau)z + R$, where $(1 - \tau)$ is the net-of-tax rate and R is virtual income.

In our analysis, we follow the earlier literature and parameterize the utility function to a quasi-linear form as follows:

$$u(c, z) = c - \frac{w^i}{1 + \frac{1}{e}} \left(\frac{z}{w^i} \right)^{1 + \frac{1}{e}}, \quad (4)$$

where w^i is an ability (productivity) parameter over which individuals are heterogeneous. Thus, the utility maximization with respect to z gives the optimal income choice for an individual, $z^* = w^i (1 - \tau)^e$, where e is the earnings elasticity parameter with respect to τ , capturing the behavioral responses to taxes. Thus, the earnings location choices of an individual i are determined in this model by innate productivity w^i , and the response to taxes determined by e . More generally, earnings elasticity in the continuous-choice model can be expressed as

$$e = \frac{dz}{d(1 - \tau)} \frac{1 - \tau}{z}. \quad (5)$$

Next, suppose that there is an increase in the tax rate above a certain income threshold. In the continuous model, if we start from individual's optimal income choice z^* and do not change the tax system applied to this location, individuals will not respond by changing their earnings. Therefore, if tax rates are changed at a higher earnings level, this model cannot capture any responses occurring to individuals below the point of the (marginal) tax rate change. Hence, the continuous model cannot capture nor explain earnings responses from below the income threshold in our empirical case.

A popular extension to the canonical model is optimization frictions that attenuate behavioral responses to tax incentives (see Chetty *et al.* (2011), Chetty (2012), Chetty *et al.* (2013), Chetty and Saez (2013), and Kleven and Waseem (2013)). Optimization frictions typically considered in the literature include job switching costs, salience of tax rules or unawareness of tax incentives that are built upon the continuous-choice model. The parameterized model could be augmented by adding a friction parameter, $a \in (0, 1)$, to the utility function. If a is close to one, responses to taxes would be minimal, and if a is close to zero, responses to taxes would occur according to the standard model. The utility function then becomes as follows

$$u(c, z) = c - \frac{w^i}{1 + \frac{1}{\epsilon(1-a)}} \left(\frac{z}{w^i} \right)^{1 + \frac{1}{\epsilon(1-a)}} \quad (6)$$

From the above equation it becomes clear that considering these types of optimization frictions merely reduce the responsiveness to taxes, but they do not alter individual earnings responses to taxes in a more fundamental manner. Therefore, this model does not produce earnings responses from below the earnings level the tax rate is changed.¹²

We can further alter the canonical framework by adding optimization errors to the model, arising from an unanticipated shock to the initially chosen income. A simple approach to include optimization errors is to consider an error parameter drawn from some distribution, $r \in f(r)$. First, a taxpayer makes an optimal earnings choice z^* , and then the optimization error alters this choice by r , so that the final outcome is $z^* - r$. A crucial aspect of optimization errors is that they are unanticipated and thus do not enter the model when the optimal earnings choice is made. Therefore, these kinds of frictions would typically cause only small deviations in income and lead to, for example, some individuals being located in the dominated range above a tax notch. However, optimization errors do not, by definition, induce large responses to changes in tax incentives, and therefore cannot explain income responses from a broad income range below a notch.

Next, we consider a model that can explain broader changes in earnings as a response

¹²Heterogeneous adjustment costs do not change this intuition, but could explain why some taxpayers respond to taxes while others do not.

to a tax rate change, namely discrete earnings choices. We define discrete earnings choices as having a non-continuous and relatively small number of alternative earnings locations from which an individual can choose from. One motivation for a discrete earnings choice set is, for example, that wage earners can typically choose between a limited number of different employers who offer jobs with a fixed monthly salary, constituting a rather sparse choice set (see e.g. Kreiner *et al.* (2015) and Saez (1999)). In addition, earnings adjustments within jobs can include discrete wage rate and working hours opportunities, leading to a non-continuous earnings choice sets. We discuss the motivation and previous literature on discrete earnings choices in more detail below in Section 4.3.

Following Saez (2002), discrete earnings could be modeled through a constraint stating that an individual chooses her earnings level from discrete earnings locations, even conditional on them being intensive margin responses, i.e. conditional on participating in the labor market. More formally, individuals choose from a discrete set of alternative earnings location choices, $u(c_{j-1}, z_{j-1})$, $u(c_j, z_j)$, $u(c_{j+1}, z_{j+1})$, but individual preferences and the underlying ability distribution are similar as before. Thus, the utility function is the same as above but indexing the discrete earnings and consumption choices with j , denoting the available earnings choices for an individual i . The discrete choice the individual chooses is determined by which location, $j - 1$ or j , yields the highest utility:

$$u(c_{j-1}, z_{j-1}) \leq u(c_j, z_j) = c_j - \frac{w^i}{1 + \frac{1}{e}} \left(\frac{z_j}{w^i} \right)^{1 + \frac{1}{e}} \quad (7)$$

The conceptual difference between this and the continuous model is that individuals now consider their utility in all potential states, z_{j-1} , z_j , z_{j+1} . Intuitively, if the tax rate changes at either z_{j-1} or z_j , this can affect the earnings location choices of an individual even if the change in the tax rate occurs further away from the current location in the earnings distribution. In other words, even tax rate changes above the current earnings location can induce earnings responses. Therefore, the discrete choice model can rationalize much larger jumps in earnings as a response to a local tax rate change compared to any of the continuous models considered above. Consequently, it also rationalizes earnings responses from a wide income range below the notch in our empirical example.

Importantly, when earnings choices are discrete, the canonical continuous-choice elasticity measure does not accurately capture the extent of the behavioral response and the welfare effects of taxes. In this case, a correct approach is to apply the mobility elasticity formula (Saez 2002; Kreiner *et al.* 2015). Contrary to the continuous model, mobility elasticity measures the earnings responses caused by individuals moving from one discrete earnings location to another as a response to a change in the tax rate. Following the lines of Kreiner *et al.* (2015) and Saez (2002), we can express the mobility elasticity with the following equation

$$\zeta_i = \frac{dY_j}{d(c_j - c_{j-1})} \frac{c_j - c_{j-1}}{Y_j} \quad (8)$$

where $c_j - c_{j-1}$ is the difference in after-tax incomes between two earnings locations j and $j - 1$, and $Y = \sum_{k=j-1}^{j+1} (z_k g_k)$, where z_k is earnings and g_k the relative mass of individuals in each earnings location. Equation (8) thus captures the change in earnings inflicted by individuals moving between different earnings locations due to changes in tax rates.

This formula captures two important features that are missing from the continuous-choice specification when considering discrete earnings responses. First, the mobility elasticity captures earnings responses from a broader income range, taking into account individuals below the location where the tax rate change occurs. Second, the standard measure for the change in the marginal tax rate does not capture the actual change in incentives for those individuals who respond to a change in the tax rate in a region below its location. Therefore, when earnings choices are discrete, the mobility elasticity concept is the correct approach to measure the welfare loss.

Finally, Figure 6 graphically illustrates the conceptual differences between the continuous and discrete choice models when a tax notch changes its location. In the canonical case in graph (a), the indifference curves are drawn such that an individual would be bunching at the original notch, and shifts her location to the right when the location of the notch is increased. By definition, individuals below the original notch are unaffected by the reform. The discrete choice model in graph (b) includes the same budget set, but the individual now faces a constraint that only certain discrete earnings locations are feasible. Under the old location of the notch, the individual would be located in the first possible earnings level below the notch. When the notch is relocated, the next discrete location above the notch becomes more attractive. The difference between the continuous and discrete models in this illustration is that the earnings response is greater in the discrete model, and occurs from a region below the original notch point. Additionally, in the discrete model, the region above the notch is no longer necessarily strictly dominated. Depending on the discrete earnings choices that are available, it is possible that the best available discrete earnings location for an individual is located just above the notch.

To summarize, this analysis illustrates that a discrete earnings choice model is the most likely candidate to explain the large and wide-ranging income responses we observe in our empirical analysis. Importantly, none of the typical optimization frictions discussed in the literature, such as inattention or optimization errors, do not consistently produce these types of responses. In Section 5.2, we utilize a simple simulation model to further illustrate this. Moreover, we demonstrate that when earnings choices are discrete, the canonical continuous model does not correctly measure the earnings elasticity and the associated welfare loss of taxes, and that the mobility elasticity estimation is the correct approach in this context. We demonstrate this feature and discuss its relevance within more general empirical applications in Section 5.3.

4.2 Mobility elasticity for students

Next, we utilize the mobility elasticity formula presented above to characterize the earnings elasticity of students. We follow the formula in equation (8) and assume a simple discrete earnings model implying that students who respond to the reform choose from two earnings locations: below the old income threshold before the reform and above it after the reform (see Figure 3). The relocation of the threshold induced a large increase in incentives to earn income above the old income threshold, thus creating a large change in the net earnings of the students who responded to the reform and relocated themselves from below to above the old threshold.

The average real earnings in 2007 and 2008 were 7,116 euros and 7,529 euros respectively, implying a 413 euro and 6% average increase in earnings within an income range of 2,000–18,000 euros. This denotes the response in earnings to the change in the location of the notch. Consistent with the evidence in Figure 4, we find no significant changes in average real earnings between other years than 2007–2008. To measure the implicit change in tax incentives and to simplify the calculations, we aggregate the choices into one average location below and one above the old income threshold. The real average gross earnings were 6,008 euros below and 11,821 euros above the old threshold in 2008. Then, we define the change in net incomes between these two earnings locations by using the tax and subsidy rules before and after the 2008 reform. Due to the shifting of the notch, the net earnings difference between these locations increased from 3,534 to 4,807 euros, or by 36%.

To calculate the mobility elasticity, we relate the average change in log gross earnings between 2007–2008 to the log change in the difference of net incomes between the two locations (below and above the old threshold). This delivers a mobility elasticity estimate of 0.184.¹³ Therefore, even though the reform caused large and distinctive earnings responses from a wide income range, the strong incentive change implies that the earnings elasticity estimate is nevertheless modest. In Section 5, we discuss the magnitude of this estimate in relation to the estimates derived using the traditional continuous-choice approach.

4.3 Further support for discrete earnings choices

Our reduced-form results combined with the theoretical considerations presented above point to the direction that the labor supply model for wage earners features discrete earnings constraints. Next, we present further arguments from earlier literature and institutional settings for why the labor supply model is likely to include discrete earnings choices. More importantly, we present sub-sample results we divide our baseline sample by

¹³Table A2 in Appendix A presents the variables used to calculate the mobility elasticity estimate.

the discreteness of labor markets the students work in, further supporting the hypothesis that earnings choices are discrete rather than continuous.

As mentioned above, the idea of discrete and inflexible labor markets is not new in the economics literature. First, non-continuous labor supply choices have been analyzed and discussed in the theoretical and structural labor supply literature. For example, Lundberg and Dickens (1993) provide a theoretical framework including a finite choice set for available working hours. Saez (1999) argues that workers with specific education, training and occupational skills rarely face a large set of available jobs that match their skills and preferences, thus limiting their potential to flexibly adjust to changes in tax rates by changing their job. Kreiner *et al.* (2015) build and estimate a structural search model with discrete choices, motivated with similar arguments as Saez (1999). Other structural labor supply models often assume that labor supply choices (working hours) are discrete in nature (see e.g. van Soest (1995) and Beffy *et al.* (2018), and Löffler *et al.* (2018) for a recent review), stemming from the observation that working hours often tend to cluster at certain focal points in the distribution. Manning (2003) discusses the role of labor market power of employers in affecting labor supply responses and working hours choices of employees.

There are a limited number of studies providing reduced-form evidence on discrete earnings choices. One exception is Blundell *et al.* (2008), who estimate the intensive margin labor supply responses of single mothers to changes in various in-work benefit programs in the UK. They find that the responses are governed by discrete working hours responses between jobs rather than continuous labor supply or wage rate adjustments.

Various institutional features provide practical reasons for why annual earnings choices could be discrete at the individual level. In many cases, wage earnings arise from hourly or monthly wages that are multiplied by time spent on working. Both of these elements, wage rates and working time, typically include discrete or discontinuous elements. First, in many types of labor markets, wage rates and working hours are regulated either by legislation or by collective agreements. These include, for example, minimum wage and minimum weekly working hours regulations. Second, discreteness arises from employment contracts between workers and firms. Such contracts are often mandated in legislation, but even in the absence of that both employees and employers can benefit from forming a contract. For a worker, contracts ensure a predictable level of future earnings, and for the employer they guarantee a sufficient labor force for a given time period. These employment contracts typically set either the wage rate or working hours or both for a fixed period of time and also include a fixed notice period, inducing discreteness in available labor supply and earnings choices for an individual.

In order to offer some stylized micro-level empirical evidence supporting the discrete individual-level earnings choices, we use wage rate and working hours register data from the Structure of Earnings Statistics provided by Statistics Finland, offering de-

tailed occupation-level data on wage rates and working hours. First, Figure 7 presents two pieces of descriptive evidence of the discreteness in wage rates at certain industries, namely transportation and cleaning services. The hourly base wage rate distribution of bus drivers in graph (a) illustrates that while there is overall variation in hourly wage rates, the distribution has clear focal points at the wage rates stemming from the collectively agreed wage schedules. Similarly, the cleaners' wage rate distribution relative to the personal minimum wage presented in graph (b) illustrates that most cleaners receive the minimum wage. Second, Graph (a) in Figure 8 shows the weekly working hours for all wage earners in Finland, highlighting the discreteness in hours worked and a very distinctive spike in full-time working hours (typically 37.5 or 36 hours per week in Finland). In graph (b) of the same figure, we exclude full-time work to underline that there are also clear focal points in the distribution of part-time work.

Finally, we provide additional evidence of how discrete earnings choices affect our empirical results. We examine the responses of two specific subgroups of students working in labor markets with more and less available discrete earnings choices: those who work in the public sector, or research, manufacturing and construction in the private sector, and those working in restaurants, bars and cafes, hotels and other accommodation services, cleaning and security services, and retail sales such as supermarkets and gas stations. The latter group has typically more flexible working hours and are subject to hourly rather than monthly wages, compared to the first group. Therefore, students in the latter group are more likely to have a larger number of available discrete earnings choices compared to the first group.

Figure 9 shows that the changes in the earnings distribution from a broader income range after the 2008 reform are smaller for those students who work in less discrete labor markets (6.14(1.71)), and larger for those working in jobs with less available discrete choices (10.94(1.10)). This difference in the amount of shifting in the distributions strongly support our hypothesis that discrete earnings responses are a key factor explaining the main result of widespread income responses far below the notch. However, we detect statistically significant responses from a wider income range also for those in the first group, illustrating that discrete earnings choices can induce relevant constraints even in more flexible labor markets, consistent with our notion that discrete earnings is a feasible framework to consider for all wage earners. Also, these findings suggest that the available discrete earnings choices can differ significantly between different types of labor markets, which can lead to differences in the estimates of observed behavioral responses to tax incentives. We discuss this in more detail in Section 5.3.

5 Simulations and broader implications of discrete earnings

In this section, we further analyze the earnings response mechanisms and present evidence for the welfare implications of discrete earnings choices. First, we present the simulation model we use in this analysis. Then, we provide simulation results that illustrate the mechanisms behind the income responses of students under different model assumptions, highlighting that the discrete earnings model as opposed to continuous-choice model qualitatively matches our empirical findings. The analysis proceeds with presenting how the standard reduced-form tax elasticity estimates based on the continuous model are biased when earnings responses are discrete. Furthermore, we illustrate that discrete earnings responses can under certain assumptions provide an intuitive explanation for why we typically observe small or modest earnings responses to taxes among wage earners. Moreover, we discuss the implications of discrete earnings choices on local bunching estimation.

5.1 Simulation model

We build the simulation model on the theoretical framework presented in Section 4. The individual utility function is given in equation (4), where the e parameter governs the disutility from earnings supply and would correspond to the elasticity with respect to taxes in the continuous model. The discrete model has the same utility function and adds a fixed number of discrete choices to the individual decision problem as an additional constraint. The budget set for individuals comes from the tax system included in the simulations which we discuss in detail below.

The model assumes an underlying ability distribution from which each individual i receives a predetermined draw w_i . This predetermined draw represents earnings in the absence of responses to the tax system. Our parameterized ability distribution is presented in Figure A2 in Appendix A.¹⁴ When the model involves a discrete earnings choice sets as constraints, they are drawn from the discrete choice distribution presented in Figure A3 in Appendix A. The number of choices drawn can be altered in different specifications, and the draws vary between different individuals. Therefore, even when the individual-level earnings choices are discrete, the overall earnings distribution is smooth.

¹⁴The distribution is a combination of power distributions and normal distributions, which gives an approximate match for the shape of the empirical earnings distribution of students in our empirical case. In general, our results are not very sensitive to different underlying ability distributions that roughly match the empirical earnings distribution of students.

5.2 Discrete vs. continuous model

First, we simulate earnings distributions when the budget set includes an income notch, a scenario that resembles our empirical case for university students. In this exercise, we assume parameters given in Table 3. The marginal income tax rate is set to 22% below the notch and a high marginal tax rate of 61% is applied above the notch, constituting a simplified linear version of the actual budget set for students including multiple notches above the income threshold (see Figure 1). The size of the notch, i.e. the size of the drop in disposable income at the income threshold, is 500 euros and the notch is relocated from 9,000 to 12,000 euros in the simulated reform.

Figure 10 shows the simulated earnings distributions within an income interval of 0–25,000 euros using the baseline continuous model and 0.2 as the value for the e parameter. The different panels in the figure present the baseline continuous model (panel a), baseline model added with adjustment frictions (panel b) and the baseline model added with both adjustment friction and optimization errors (panel c).¹⁵

First, we find that bunching at the income threshold is sharp and sizable both before and after the reform in panel (a). Adding adjustment costs to the model according to equation (6) in panel (b) only leads to slightly lower bunching and slightly more individuals being located just above the notch, but no changes in earnings from a broader income range. Adding i.i.d. earnings shocks as optimization errors in panel (c) yields more diffuse bunching, but again no earnings responses from a wider income range below the old income threshold.¹⁶

Next, we add discrete earnings choices as an additional constraint. Figure 11 illustrates the earnings distributions using 30, 15, 10 and 5 available earnings choices for each individual and assuming an underlying e parameter of 0.5. With a discrete choice constraint included in the model, the earnings distributions and the response to the reform qualitatively resemble the empirical distributions (see Figure 3). In particular, when the number of discrete choices is set to 10 (on average, earnings jumps of 2,500 euros or 10–30% relative to the baseline income level), the qualitative shape of the distributions largely resemble their empirical counterparts. First, the earnings distribution shifts to the right from a relatively wide income range below the old threshold. Second, the shape and amount of excess bunching below the threshold are approximately of the same order of magnitude as in our empirical case. Note that scattered local bunching below the threshold only results from including discrete earnings choices as a constraint, as this

¹⁵We assume heterogeneous adjustment frictions represented by a uniformly distributed parameter a in the unit interval. Each individual has a different and independent draw from this distribution. The earnings shocks related to optimization errors are normal distributed mean-zero income shocks with a standard deviation of 800 euros.

¹⁶If we were to assume only negative income shocks we would get diffuse bunching only below the notch, similarly as in the empirical distribution. However, such asymmetric shocks cannot be not easily justified.

model does not include optimization errors or adjustment frictions.

Overall, these findings further support the conclusion in Section 4.1 that the canonical continuous-choice model cannot explain our empirical findings, not even when complemented with adjustment frictions or optimization errors. In contrast, discrete earnings constraints are able to at least qualitatively explain the earnings response from far below the old income threshold.

5.3 Tax elasticity estimation

In this section, we show the impact of the earnings supply model on earnings response estimates and the potential biases therein based on the theoretical considerations presented in Section 4.1. Our starting point is that most of the existing reduced-form literature on earnings responses to taxes, mainly consisting of various elasticity of taxable income (ETI) estimates, assume some version of a standard continuous-choice model, either explicitly or implicitly (Saez *et al.* 2012; Kreiner *et al.* 2015). We analyze how these elasticity estimates depend on the underlying modeling assumptions, focusing on the implications when the underlying earnings adjustment process is discrete rather than continuous.

We simulate earnings distributions under a kinked tax schedule where the marginal tax rate is increased above a predetermined income level. This resembles a typical setup for analyzing earnings responses to tax rate changes in the empirical literature. In our baseline case, the tax schedule features a basic marginal income tax rate of 40% which jumps to 60% from earnings of 10,000 euros onward. A reform repeals this kink and reduces the tax rate to 40% also from the 10,000 euro threshold onward. As above, we model individual behavior within an income range of 0–25,000 euros with an average income being approximately 6,000 euros. Thus, this set-up would correspond to a large marginal tax rate cut for higher incomes, in a similar vein as the widely studied Tax Reform Act of 1986 in the US (Feldstein 1999, Saez *et al.* 2012).

Figure 12 illustrates the simulated earnings distributions before and after the reform for our baseline model assuming an e parameter of 0.5 and 10 available discrete earnings choices for each individual. For a comparison, Figure 13 presents the distributions assuming continuous earnings. These figures highlight similar differences between the models as our previous analysis of the study subsidy notch. First, there is significant and sharp local bunching in the continuous case but no clear bunching at the kink when earnings choices are discrete. Second, the earnings distribution changes shape from a broader region below the kink point in the discrete case, similarly as in our empirical example. In contrast, in the continuous case the responses are limited to very sharp local bunching and earnings responses of individuals above the kink point, stemming from the assumptions of the continuous-choice model.

Next, we analyze in detail how the different earnings choice mechanisms affect the

estimated elasticities. Typically, panel data estimates for ETI are conducted using a differences-in-differences type of an approach (see e.g. Saez *et al.* (2012) for more details on the typical estimation procedure). To obtain an elasticity estimate, the earnings development of individuals who face a change in their marginal tax rate based on their pre-reform earnings is related to similar individuals who do not face a change in their taxes. Therefore, this approach ignores any responses stemming from below the original kink by definition. Also, the tax rate change in this model is only measured for those individuals whose pre-reform earnings are above the original tax rate kink. Formally, the standard estimation is usually performed by regressing the change in log income on the change in the log net-of-tax rate as follows

$$d \log z_{it} = \varepsilon d \log(1 - \tau)_{it} + \epsilon_{it} \quad (9)$$

where subscript i refers to the individual and t time, z are (taxable) earnings, τ is the marginal tax rate, ϵ denotes the error term, and ε the estimated average earnings elasticity with respect to the net-of-tax rate.

The empirical mobility elasticity formula we use in the analysis differs from this specification in two important ways. First, we account for the fact that when earnings choices are discrete, also individuals below the original kink point can be affected by the reform. We denote the tax rate (m) as the difference between the net earnings relative to gross earnings in the two earnings locations i and $i - 1$, as discussed in Section 4.1. Empirically, m can be calculated by estimating the average change in income due to the reform by income bin and denoting the initial earnings location as $i - 1$ and the new location (pre-reform income plus the average change in income) as i . This implies that the change in the tax rate is potentially non-zero for a large part of the earnings distribution also below the kink. Formally, the mobility elasticity is estimated with the following equation

$$d \log z_{it} = \zeta d \log(1 - m)_{it} + \epsilon_{it} \quad (10)$$

where ζ denotes the average mobility elasticity.

Figure 14 presents the earnings responses and tax rate changes associated with the reform in both the discrete and continuous cases. In the figure, the responses are denoted in terms of pre-reform earnings (horizontal axis) using 200 euro income bins. As can be expected, the change in log income is zero for those originally below the kink in the continuous model. The earnings response is a constant 0.2 above the narrow bunching window just above the original kink point. Following the assumptions in the continuous model, the change in the log of net-of-tax rate is 0.4 above the kink, and zero below it.

These patterns are different in the discrete case when we assume 10 available earnings locations for each individual. First, the change in log earnings is not zero in the region below the original kink. This stems from the fact that in that range there are a number

of individuals whose next available earnings location is in the region where the tax rate change has occurred. Thus, the reform can affect the optimal earnings location choices also below the kink point, and after the reform individuals originally below the kink relocated themselves into the income range above the kink where the tax rate was reduced from 60 to 40%.

Second, the average earnings response is smaller than in the continuous model in the region above the kink. This stems from the fact that the discrete choice set limits the earnings responses for many individuals, and therefore, only a smaller share of individuals change their earnings location after the reform. In other words, even though the individual-level jumps in earnings are in this case relatively large (10–30% relative to the baseline income level), not all individuals respond to the reform by changing their location. In contrast, the continuous model implicitly assumes that all individuals above the kink respond to the reform with a similar earnings response.

Third, the change in the tax rate m reduces the further we move below the original kink point, because the further we move away from the kink the greater is the distance between the two earnings locations i and $i - 1$ we are considering, and thus the relative change in net income due to the tax change becomes smaller. The change in net of average tax rate is equal to the change in net of marginal tax rate above the kink, because then we are comparing the relative change in net income caused by the same tax schedule between the two locations.

This discrepancy between the models has a significant effect on the implied tax elasticity estimate. Figure 15 shows what the corresponding tax elasticity estimates would be in the continuous and discrete models. Unsurprisingly, the continuous model produces a constant elasticity estimate of 0.5 above the kink point, which equals the chosen e parameter in the model. The mobility elasticity in the discrete model, however, is not constant across the distribution and is smaller than in the continuous case. The average mobility elasticity in this example is 0.14, and the average elasticity is 0.2 above the kink and 0.12 below it. These numbers imply that when earnings choices are discrete, the true elasticity estimate can be significantly different than what we would predict based on the canonical continuous model. Therefore, as it is likely that the discrete earnings constraint induces a relevant factor affecting the earnings responses of workers in many types of labor markets, it provides an intuitive explanation for why we typically observe relatively small earnings responses to tax rate changes among wage earners.

Moreover, earnings responses under discrete choices can produce a notable bias to the typical differences-in-differences estimator used in the empirical ETI literature. In reduced-form applications using actual tax reforms and administrative data, we typically cannot relate the earnings development of those who faced a tax rate change to individuals with the exact same pre-reform earnings levels who did not face a change in taxes, as we did when we derived the simple tax elasticity estimates above. In contrast, we

typically use the earnings of (more or less) similar individuals who did not face a change in their marginal tax rate based on their pre-reform earnings as a control group for those who did, as in a seminal paper by Gruber and Saez (2002). Next, we characterize this “naive” differences-in-differences ETI estimate using the individuals above the kink as the treatment group, and individuals within an income range of 8,000 euros below the kink as the control group. Following the intuition in the continuous-choice model, individuals below the kink are assumed to be unaffected by the reform, thus constituting a control group comprising of approximately similar individuals than those above the kink who face a tax rate cut based on their pre-reform earnings. However, the identifying assumptions for the control group are clearly violated when earnings choices are discrete, as also individuals below the kink point are affected by the reform, as already illustrated above.¹⁷

Table 4 collects the elasticity estimates in the discrete case with 10 available earnings choices and varying the location of the original kink point. The table presents the naive ETI estimates described above and the mobility elasticity estimates which are estimated for the whole data utilizing a pure control group of individuals in the same income range who did not face the tax reform. The results show that the biased naive ETI estimate is always lower than the mobility elasticity estimate in the simulated reform. Intuitively, this stems from the fact that the simulated reform also induced an average earnings increase in the control group below the kink, therefore reducing the naive differences-in-differences estimate. We further illustrate the source of this bias in Figure 16. In the figure, we show the estimates for the mobility elasticity both below and above the kink point as in Figure 15 above. The naive approach would utilize the group below the kink as a control group and those above the kink as the treatment group. Therefore, the associated differences-in-differences estimator would simply deduct the responses of the control group (elasticity estimate 0.12) from the treatment group (0.20), producing a small and downward-biased elasticity estimate of 0.08. In contrast, no such bias occurs if the earnings choices would be continuous, as in this case the response is limited to those above the kink, and the estimate (0.5) is therefore not affected by the earnings responses of the control group below the kink.

Furthermore, a comparison between the naive ETI and mobility elasticity estimates in Table 4 suggests that the size of this bias can be notable. For example, using the baseline 10,000 euro kink point, the mobility elasticity estimate is almost twice as large as the naive ETI estimate (0.177 vs. 0.098). More broadly, this downward bias in the ETI estimate provides a further explanation for why the observed earnings elasticity estimates in the literature are small or modest, typically ranging from 0–0.5 for wage earners (see Saez *et al.* (2012) and Neisser (2019) for a recent meta-analysis).

¹⁷Also, there are other relevant issues that need be taken into account when estimating ETI using actual tax reforms, such as the endogeneity of individual-level tax rate changes, mean reversion of income and potential differences in income growth rates at different income levels (Saez *et al.* 2012). However, as we use a simulated reform and simulated data, we do not need to account for such issues in this analysis.

Of course, the extent of the bias and the size of the mobility elasticity estimates depend on the assumptions and choices made in the simulation. Furthermore, this variation enables us to illustrate the circumstances in which the bias is smaller or larger, and how different details of the simulated tax reform affect the size of the mobility elasticity estimate. Tables A3–A5 in Appendix A present the estimates when varying the number of available set of earnings choices and the location of the original kink point. First, we observe that at least in this simulation framework, the higher the original kink is in the income distribution the lower is the mobility elasticity estimate. This stems from the fact that tax rate changes at lower income levels induce a larger number of individuals to change their earnings locations as a response to the reform. As the number of available earnings choices is smaller at larger income levels in our simulations, there is a smaller potential to respond to tax rate changes at the upper tail of the distribution. Nevertheless, the bias between the naive ETI estimate and the mobility estimate is economically significant irrespective of the location of the original kink point.

Second, as can be expected, the bias in the naive ETI estimate compared to the mobility estimate is smaller the less discrete the earnings choices are. This is due to the fact that responses below the original kink are less substantial when earnings choices are more flexible. Therefore, in the cases where individuals can easily adjust their earnings, such as for business owners and top income earners, the earnings elasticity estimates are likely to be less biased, and at the same time, larger than for regular wage earners who are more likely to face a discrete earnings choice set. This is also what is typically observed in the literature (Saez *et al.* 2012). Relatedly, the mobility elasticity estimates themselves are also higher the larger is the number of available earnings choices. Moreover, the naive ETI estimates can even have negative values when the original kink is at a lower income level and the responses from below the original kink are more pronounced, thus creating a larger bias in the differences-in-differences estimates.

Overall, our findings from these simulations imply that the interpretation of the differences-in-differences ETI estimates in earlier studies can be very sensitive to the underlying labor market constraints and the earnings adjustment mechanisms. Nevertheless, this bias is not relevant for all quasi-experimental tax elasticity estimates. For example, the literature studying the impact of the earned income tax credit (EITC) on earnings typically applies an estimation strategy where women without children are used as a control group for women with children who are eligible for larger tax credits (see Kleven (2019) for a recent review). In these types of set-ups, the control group is unaffected by the changes in taxes occurring for the treatment group, implying that potential discrete earnings responses do not bias the estimator in this case.

Finally, the discrete earnings constraint and its impact on the estimated earnings elasticity have implications for the welfare analysis of taxation. The fundamental aim of this analysis is to estimate a structural long-run elasticity unaffected by any short-run

frictions (Chetty 2012; Kleven and Waseem 2013). In contrast to frictions that can be thought to be more temporary, such as inattention to taxes or optimization errors, discrete earnings choices are likely to constitute a more permanent friction, as argued above in Section 4.3. Therefore, the observed elasticities under discrete earnings choices are likely to be closer to the underlying structural elasticity. Therefore, discrete earnings provide an intuitive explanation for why the observed moderate earnings elasticity estimates for wage earners can also represent a sufficient statistic for welfare analysis.

5.4 Bunching estimation

In addition, the local bunching model relies on the continuous-choice framework (Saez 2010; Kleven and Waseem 2013). However, the discrete earnings constraint sets some limitations to the bunching method and the interpretation of the observed local bunching response. Given that our empirical analysis documented that earnings responses can be discrete even in the part-time labor market, these limitations might be important and hinder the capability of the bunching method in delivering relevant parameters for welfare analysis in the context of labor market outcomes.

The limitations of the bunching method can be illustrated using Figure 17, which presents the simulated earnings distributions using different available discrete choice sets and the hypothetical tax rate kink reform analyzed above. The figure shows that under the discrete choice constraint the relocation choices of individuals can occur far below the local discontinuity, even when assuming as much as 30 available earnings choices within an income range of 0–25,000 euros. Therefore, local bunching responses to a discontinuity in the budget set have a limited potential to capture all relevant responses to tax incentives when earnings choices are discrete.

Furthermore, discrete earnings imply that the surrounding distribution outside the local bunching region just around the discontinuity cannot be necessarily used to estimate a credible counterfactual describing the shape of the distribution in the absence of a notch or a kink (see Appendix B for a detailed discussion on local bunching estimation). If the shape of the distribution further away from the discontinuity is also affected, we can no longer rely on the idea that the surrounding density provides us a credible counterfactual unaffected by behavioral responses. Also, without a set-up where the location of a notch or kink is relocated, we have no obvious approach to evaluate how the surrounding density would be affected. Combining the difficulty of measuring a reliable counterfactual with the notion that local bunching responses cannot capture all intensive margin behavioral responses under discrete earnings choices limit the capability of the bunching approach in delivering reliable estimates for the behavioral earnings responses to local tax rate changes.¹⁸

¹⁸Furthermore, the region of dominated choice just above an income notch is not necessarily a sub-

Relatedly, the elasticity estimate derived using the bunching method cannot be used to calculate the relationship between the change in the tax rate and the local excess mass when earnings choices are discrete. As in the case of the standard ETI estimation, both the nominator (earnings response) and the denominator (change in the tax rate) of the elasticity formula are incorrectly specified and measured. A more coherent strategy is to utilize the mobility elasticity estimate discussed in Section 4.2, linking the fraction of individuals moving to their next available location to their implicit change in tax incentives.

How does our mobility elasticity estimate (0.18) compare to the estimate derived from the local bunching model? Using the reduced-form earnings elasticity formula for income notches presented in Kleven and Waseem (2013), we obtain a local elasticity estimate of 0.065(0.007) for students with 9 subsidy months (see Appendix B for more details on the estimation procedure). Therefore, in this case, the local continuous-choice framework delivers an earnings elasticity estimate that is approximately 2.5 times smaller than in the discrete choice framework, stemming mainly from the fact that the local approach underestimates the extent of the overall behavioral response to the notch. Another approach to measure this bias is to relate the broader changes in the distribution to the local estimate derived from the continuous model. In Figure 3, we find that the broader changes in the earnings distribution are approximately 3.3 times larger than the local bunching estimate (9.81 vs. 2.93). Overall, these findings imply that when earnings choices are discrete, the traditional bunching method significantly underestimates the true earnings responses.

Consequently, discrete earnings choices can feasibly explain why we tend not to observe bunching at tax rate kinks but tend to find significant ETI estimates from the same countries and contexts. For example, in the US, bunching responses at income tax rate kinks for other individuals than the self-employed are very small or zero (Saez 2010), but the differences-in-differences estimates for ETI are typically significantly larger even for regular wage earners (see Saez *et al.* (2012) and Weber (2011)). Similar evidence is also available for Sweden (see Bastani and Selin (2014) and Blomquist and Selin (2010)) and Denmark (see Kleven and Schultz (2014) and Chetty *et al.* (2011)).

Finally, even though the local bunching approach has clear limitations when earnings choices are discrete, it can be a useful method in other contexts. For example, bunching estimates can deliver relevant evidence of behavioral responses among business owners and the self-employed or when detecting tax avoidance and evasion responses. In these instances, discrete choices is not a likely constraint for earnings responses.

optimal choice for an individual with a discrete earnings constraint. An earnings choice might be optimal even within the dominated range if other earnings choices are sufficiently far away from this region. This holds even when there are no other types of frictions such as optimization errors or inattention. Therefore, under discrete earnings choices, following the approach in Kleven and Waseem (2013) and relating the share of individuals in the dominated range to the estimated local counterfactual does not necessarily deliver us a robust measure of other frictions affecting local responses to taxes.

6 Concluding remarks

In this paper, we find clear reduced-form evidence of significant discrete earnings responses to changes in tax incentives among Finnish university students. We develop theoretical arguments and present illustrative simulation results showing that these findings cannot be explained with optimization frictions that are typically discussed in the literature, such as inattention, salience or optimization errors. Our analysis reveals that wage earners even in the part-time and temporary labor markets can face significant restrictions in their available earnings choices. This evidence together with additional descriptive findings from a more general labor market contexts highlight that discrete earnings choices are likely to be relevant in various context when analyzing labor supply and earnings responses.

Moreover, discrete earnings choices have important broader implications. Using our simulation framework, we find that the standard estimation procedure based on the canonical continuous labor supply model produces a significant negative bias to the tax elasticity estimate whenever individual earnings choices are sufficiently discrete. In addition, we find that average earnings responses are typically smaller when earnings choices are discrete, therefore suggesting that discrete earnings constraints provide a feasible explanation for why we tend to observe small or modest earnings responses to taxes among wage earners. Overall, these findings imply that the interpretation of the reduced-form estimates in earlier studies can be very sensitive to the underlying labor market constraints and the earnings adjustment mechanisms.

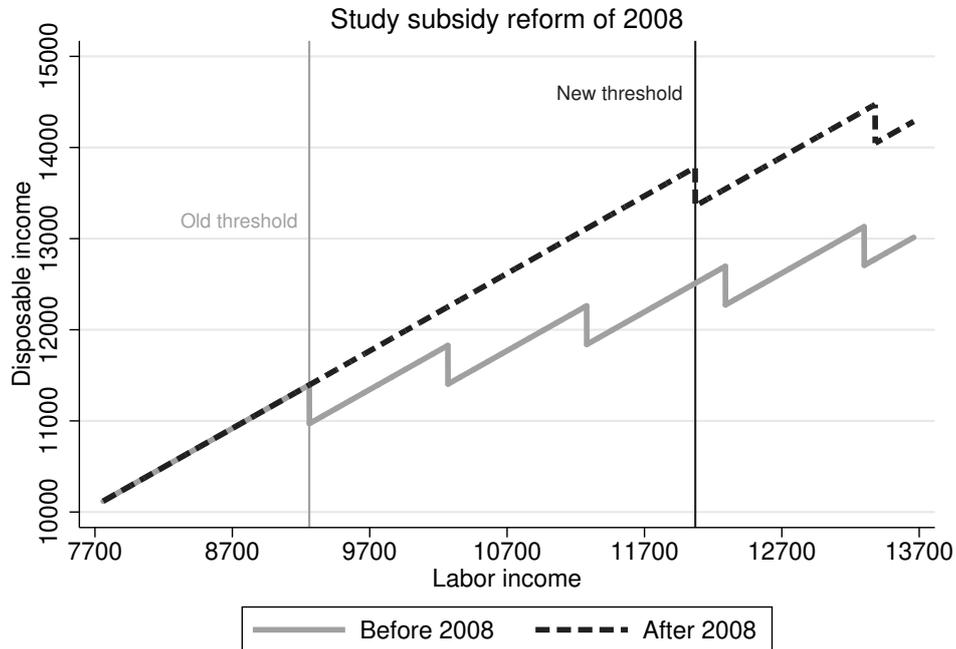
References

- [1] Bastani, S. and H. Selin (2014). Bunching and Non-bunching at Kink Points of the Swedish Tax Schedule. *Journal of Public Economics*, 109: 36–49.
- [2] Beffy, M., R. Blundell, A. Bozio, G. Laroque and M. To. (2019). Labour Supply and Taxation with Restricted Choices. *Journal of Econometrics*, 211: 16–46.
- [3] Blomquist, S., W. Newey, A. Kumar and C-Y. Liang (2018). Identifying the Effect of Taxes on Taxable Income. Working Paper.
- [4] Blomquist, S. and H. Selin (2010). Hourly Wage Rate and Taxable Labor Income Responsiveness to Changes in Marginal Tax Rates. *Journal of Public Economics*, 94(11–12): 878–889.
- [5] Blundell, R., M. Brewer and M. Francesconi (2008). Job Changes and Hours Changes: Understanding the Path of Labor Supply Adjustment. *Journal of Labor Economics*, 26(3): 421–453.

- [6] Chetty, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics*, Vol. 1, 2009.
- [7] Chetty, R. (2012). Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, 80(3): 969–1018.
- [8] Chetty R., A. Looney and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4): 1145–1177.
- [9] Chetty, R., J. Friedman, L. Pistaferri and T. Olsen (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *Quarterly Journal of Economics*, 126(2): 749–804.
- [10] Chetty, R., J. Friedman and E. Saez (2013). Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings. *American Economic Review*, 103(7): 2683–2721.
- [11] Chetty, R. and E. Saez (2013). Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients. *American Economic Journal: Applied Economics*, 5(1): 1–31.
- [12] Dickens, W., and S. Lundberg (1993). Hours Restrictions and Labor Supply. *International Economic Review*, 34 (1), 169–192.
- [13] Feldstein, M. (1999). Tax Avoidance and the Deadweight Loss of the Income Tax. *Review of Economics and Statistics*, 81(4): 674–680.
- [14] Gelber, A., D. Jones and D.W. Sacks (2019). Estimating Earnings Adjustment Frictions: Method and Evidence from the Earnings Test. *American Economic Journal: Applied Economics*, forthcoming.
- [15] Gelber, A., D. Jones, D.W. Sacks and J. Song (2018). Using Kinked Budget Sets to Estimate Extensive Margin Responses: Method and Evidence from the Social Security Earnings Test. Working Paper.
- [16] Gruber, J. and E. Saez (2002). The elasticity of taxable income: evidence and implications. *Journal of Public Economics* 84, 1–32.
- [17] Kleven, H. (2016). Bunching. *Annual Review of Economics*, Vol. 8, 2016.
- [18] Kleven, H. (2019). The EITC and the Extensive Margin: A Reappraisal. NBER Working Paper No. 26405.
- [19] Kleven, H. and E. Schultz (2014). Estimating Taxable Income Responses Using Danish Tax Reforms. *American Economic Journal: Economic Policy*, 6(4): 271–301.

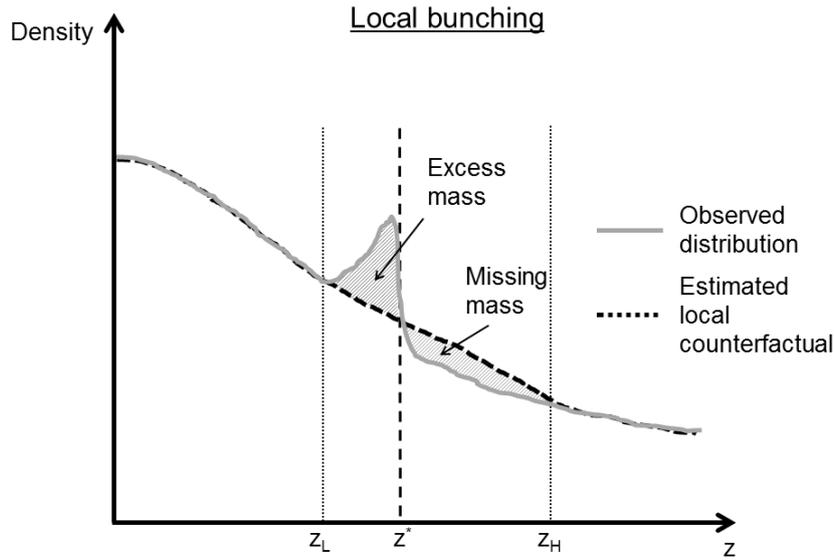
- [20] Kleven, H. and M. Waseem (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *Quarterly Journal of Economics*, 128: 699–723.
- [21] Löffler, M., A. Peichl and S. Siegloch (2018). The Sensitivity of Structural Labor Supply Estimations to Modeling Assumptions. *ifo Working Papers* 259/2018.
- [22] Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.
- [23] Neisser, C. (2019). The Elasticity of Taxable Income: A Meta-Regression Analysis. Working Paper.
- [24] Saez, E. (1999). Do Taxpayers Bunch at Kink Points? *NBER Working Papers* 7366.
- [25] Saez, E. (2002). Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses, *Quarterly Journal of Economics*, 117: 1039–1073.
- [26] Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3): 180–212.
- [27] Saez, E., J. Slemrod and S. Giertz (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, 50(1): 3–50.
- [28] Slemrod, J. (1998). Methodological Issues in Measuring and Interpreting Taxable Income Elasticities. *National Tax Journal* , 51(4): 773–788.
- [29] Søggaard, J. (2019). Labor Supply and Optimization Frictions: Evidence from the Danish Student Labor Market. *Journal of Public Economics*, 173: 125–138.
- [30] van Soest, A. (1995). Structural Models of Family Labor Supply - a Discrete Choice Approach. *Journal of Human Resources*, 30(1): 63–88.
- [31] Weber, C. (2014). Toward Obtaining a Consistent Estimate of the Elasticity of Taxable Income Using Difference-in-differences. *Journal of Public Economics*, 117: 90–103.

Figures

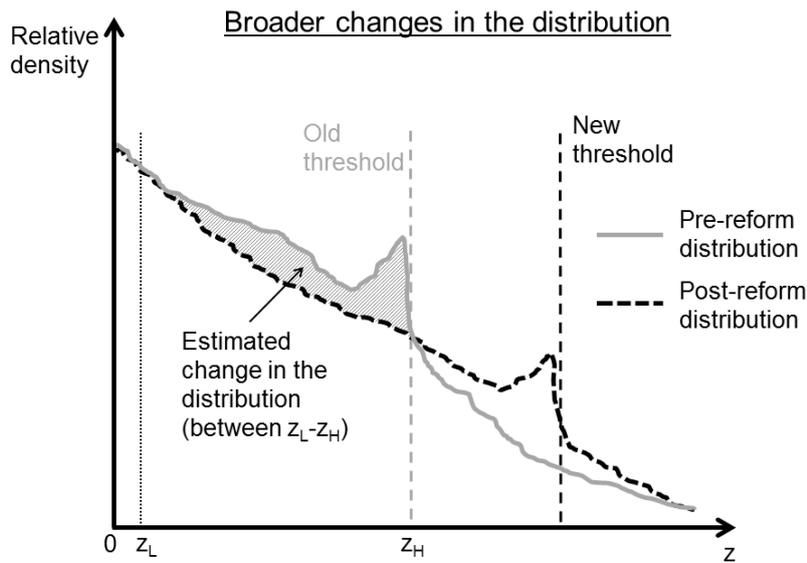


Notes: Figure presents the study subsidy schedule before (gray solid line) and after 2008 (black dashed line) for a student who collects the default 9 subsidy months. The vertical axis denotes disposable income, and horizontal axis labor income. The vertical lines denote the thresholds before (9,200 euros) and after (12,070 euros) the 2008 reform. Above the income threshold, one month of the study subsidy is reclaimed, resulting in a discontinuous drop in disposable income. Furthermore, an additional month of the subsidy is reclaimed after an additional 1,010 and 1,310 euros above the threshold before and after 2008, respectively. The figure illustrates the distinctive change in incentives caused by the increase in the income threshold in 2008, highlighting that the reform encouraged to increasing earnings above the old income threshold.

Figure 1: Disposable income at different earnings levels for students with 9 subsidy months in 2007 and 2008



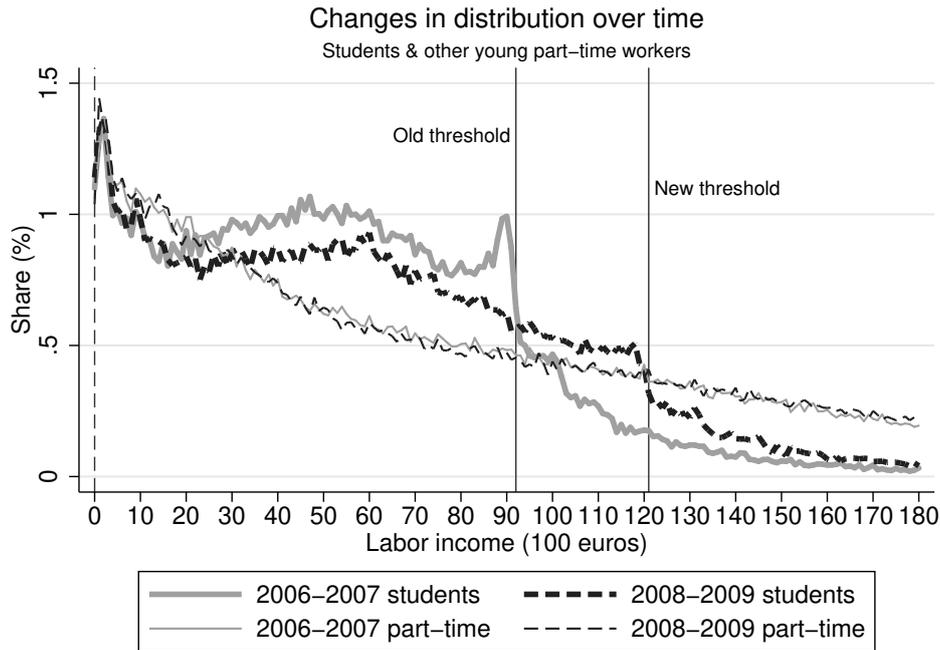
(a) Local bunching



(b) Broader changes in the distribution

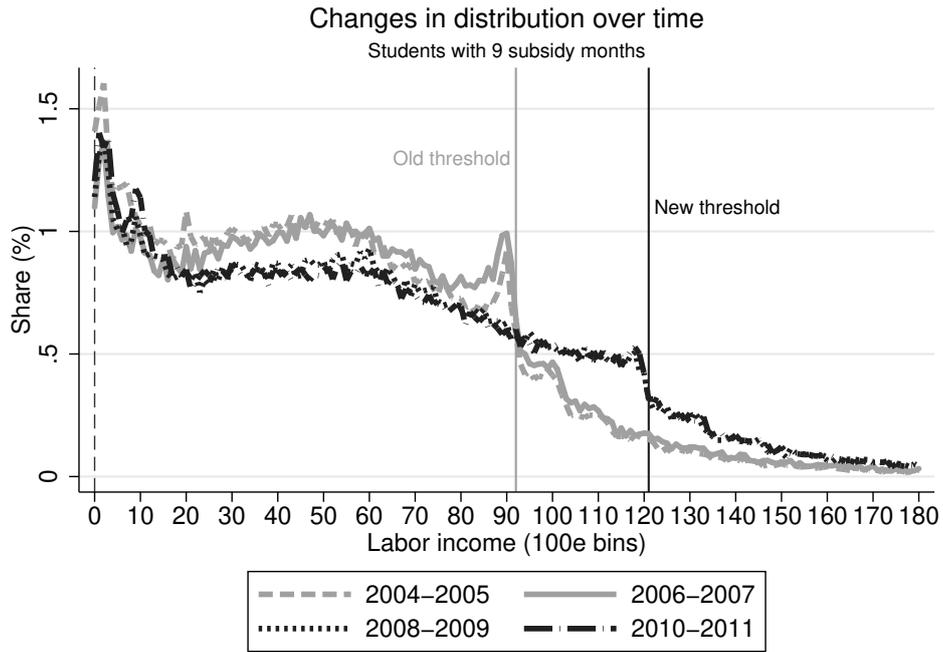
Notes: Graph (a) illustrates the local excess bunching at the income threshold in a hypothetical earnings (z) distribution (gray solid line), compared to an estimated counterfactual distribution in the absence of the threshold (black dashed line). In the figure, the threshold is denoted by z^* , and z_L and z_H denote the lower and upper limits of the bunching region. The procedure for estimating local excess bunching is described in more detail in Appendix B. Graph (b) illustrates broader changes in a hypothetical earnings distribution after an increase in the location of the threshold. The pre-reform distribution is marked with a gray solid line and the post-reform distribution with a black dashed line. z_L and z_H denote the lower and upper limits of the estimation region.

Figure 2: Local bunching and broader changes in the earnings distribution



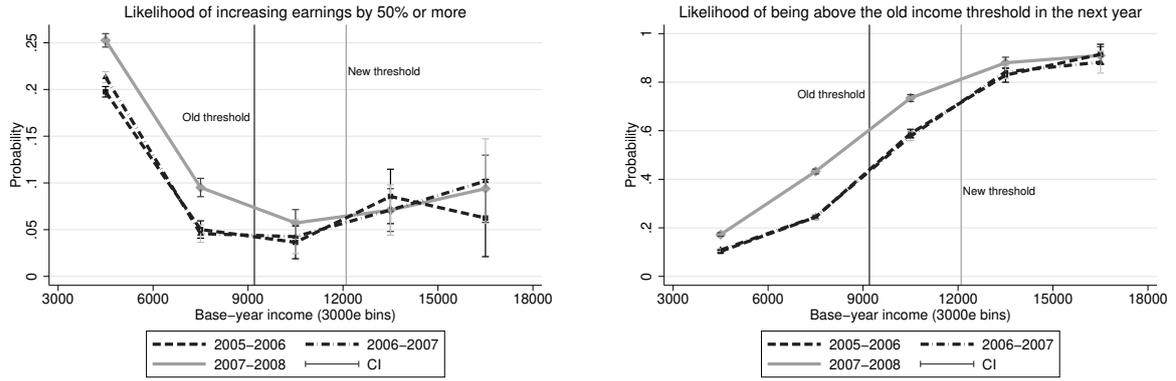
Notes: Figure presents the observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 100 euros for students with the default 9 subsidy months in each year, and for young part-time workers who are not students (see Table 2). The first vertical line at 0 denotes the lower limit in the estimation of broader earnings changes in the distribution estimated using equation (2), and the second and third lines denote the pre- and post-reform income thresholds, respectively. The figure illustrates that the earnings distribution after 2008 has a significantly different shape than before the reform, implying that the income threshold affects the shape of the whole labor earnings distribution, not just the region close to the notch point. The difference-in-differences estimate for broader changes in the distribution within an income range of 0–9,200 euros is 9.81 (standard error 1.01). The estimate for broader changes among only the student population is 10.97 (standard error 1.85), estimated using equation (1). Local bunching estimates at the threshold are 2.93(0.88) before and 1.71(0.88) after 2008, respectively. A lower limit of 1,100 euros below the threshold is used in the estimation of local bunching both before and after 2008. See Appendix B for a more detailed analysis of local bunching responses.

Figure 3: Earnings distributions of students and non-student part-time workers before and after the 2008 reform



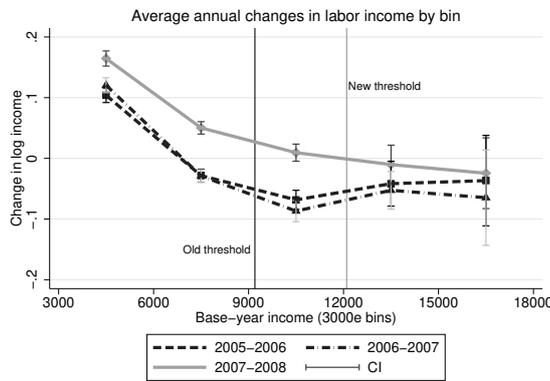
Notes: Figure presents the earnings distributions of students with 9 subsidy months in 2004–2005 (gray dashed line), 2006–2007 (gray solid line), 2008–2009 (black solid line) and 2010–2011 (black dotted line) within an income range of 0–18,000 euros in bins of 100 euros. The figure shows that the response of students occurred exactly at the time of the 2008 reform, and is not caused by gradual changes in the shape of the earnings distribution over time.

Figure 4: Earnings distributions of students in 2004–2011



(a) More than 50% earnings increases

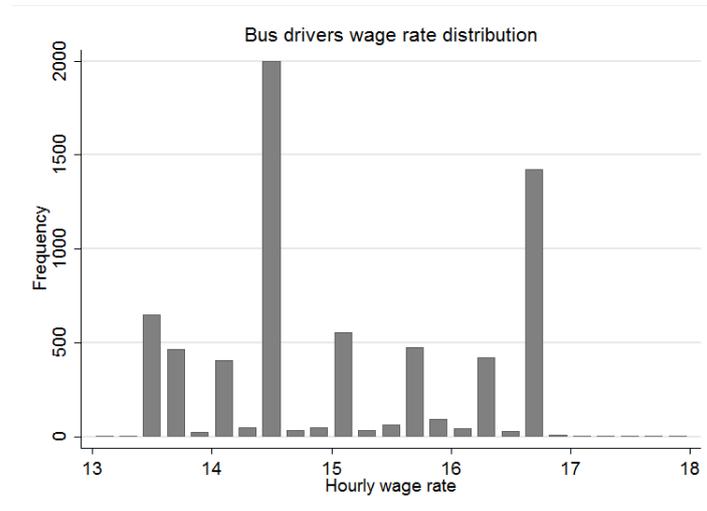
(b) Locating above the old income threshold



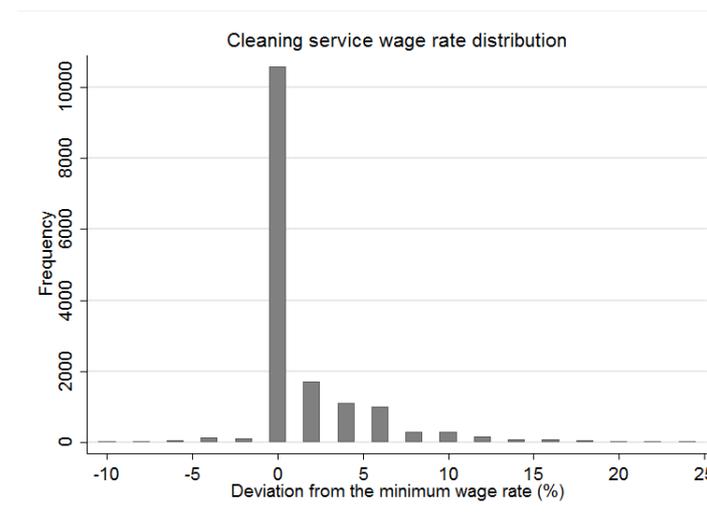
(c) Average changes in labor income

Notes: Graph (a) presents the average likelihood and 95% standard errors for increasing labor income by 50% or more relative to base-year income in base-year bins of 3,000 euros for students with 9 subsidy months. Gray solid line presents the years 2007–2008, and black dashed lines the pre-reform years 2005–2006 and 2006–2007. The graph illustrates that the likelihood of large income increases is significantly higher below the old threshold at the time of the reform compared to previous years, but there are no significant changes above the old threshold between the years. Graph (b) presents the average likelihood and 95% standard errors for locating above the old income threshold in the next year in base-year income bins of 3,000 euros. The graph shows that this likelihood increased significantly in bins below the new threshold, but there are no significant changes between the years at larger income levels. Graph (c) presents the relative average individual-level changes in real labor income (relative to 2007 real price index) with 95% standard errors in base-year bins of 3,000 euros. The graph shows that earnings increases are more prevalent below the new threshold at the time of time reform compared to previous years, but there are no significant differences above the new income threshold. Overall, these findings support the view that students responded to the relocation of the notch with large earnings increases instead of marginal earnings adjustments along the whole distribution.

Figure 5: Further evidence of discrete earnings responses to the 2008 reform



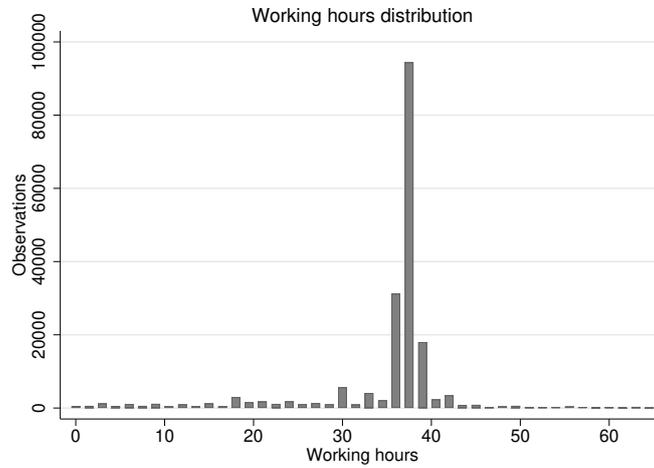
(a) Hourly wage rate distribution of bus drivers



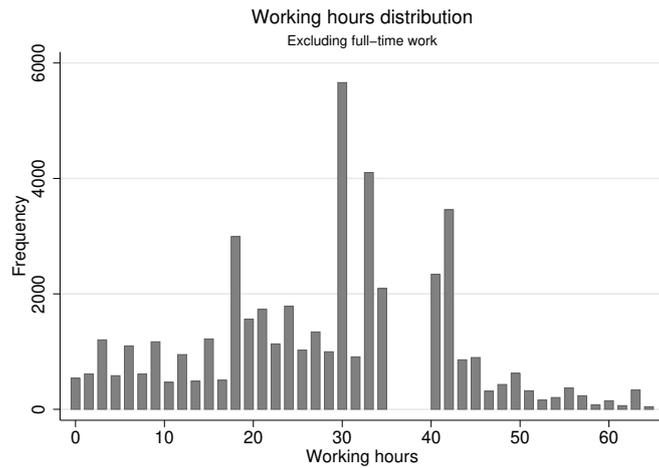
(b) Wage distribution relative to personal minimum wage, cleaning services

Notes: Figure presents general description of discrete wage rate choices of individuals using the Structure of Earnings Statistics for the year 2016 provided by Statistics Finland. Graph (a) illustrates the wage rate distribution of bus drivers. While there is overall variation in hourly wage rates, the distribution has clear focal points at the wage rates stemming from the collective agreements between the representatives of labor and employers' organizations. Therefore, from the individual point of view, wage rate changes often occur in a discontinuous manner. Graph (b) presents the wage rate distribution of individuals working in cleaning services relative to the industry-specific minimum wage, showing that a bulk of individuals in that industry are restricted by the minimum wage, indicating that continuous wage rate adjustments are rather restricted due to this regulation.

Figure 7: Discrete wage rate opportunities for wage earners



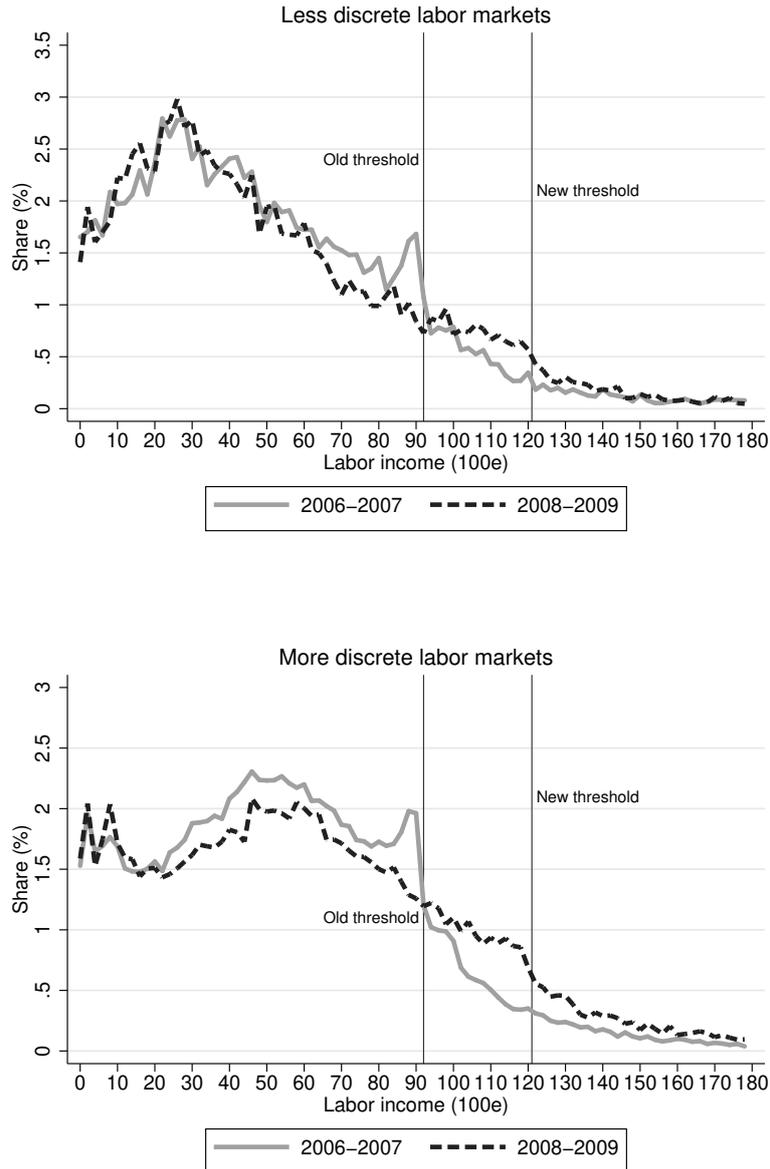
(a) Aggregate working hours distribution



(b) Excluding typical full-time working hours

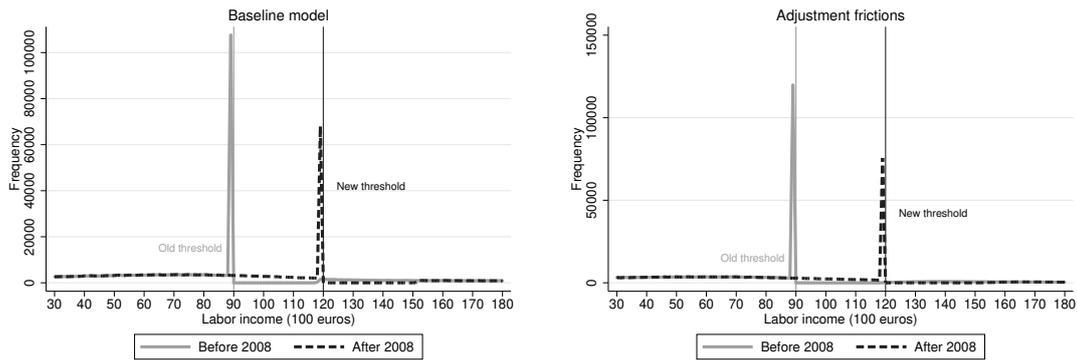
Notes: Figure presents general description of working hours choices of individuals using the Structure of Earnings Statistics for the year 2016 provided by Statistics Finland. Graph (a) presents the weekly working hours distribution for all workers. The graph illustrates that in many cases employment contracts commit workers for a full-time job for a set time period, which can be seen as a large spike in typical full-time working hours in Finland, such as 36 and 37.5 hours per week. Graph (b) shows the working hours distribution excluding the typical full-time working hours from the distribution, showing that working hours for part-time employment also tend to cluster at certain focal options, such as 30 and 18 hours per week.

Figure 8: Aggregate working hours distributions for wage earners



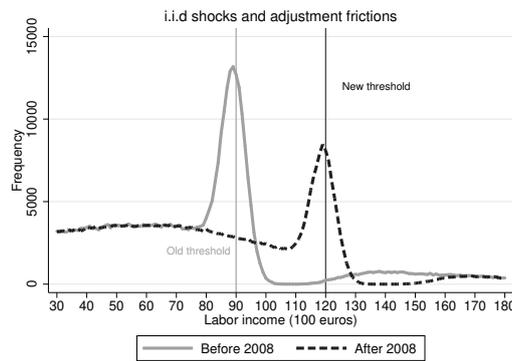
Notes: Figure presents the observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 200 euros for students with the default 9 subsidy months in each year working in different types of jobs. Jobs are categorized using firm-level industry classification codes. Less discrete labor markets include restaurants, bars and cafes, cleaning and security services, and retail sales such as supermarkets and gas stations. More discrete labor markets include public sector, and research, manufacturing and construction in the private sector. Using equation (1), the estimate for broader changes in the distribution within an income range of 0–9,200 euros for the less discrete group is 6.14(1.71), and for the more discrete group 10.94(1.10), illustrating that broader changes in the distribution are significantly more prevalent for the latter group compared to the first group.

Figure 9: Labor income distributions before and after 2008 for students working in less discrete and more discrete labor markets



(a) Baseline simulation model

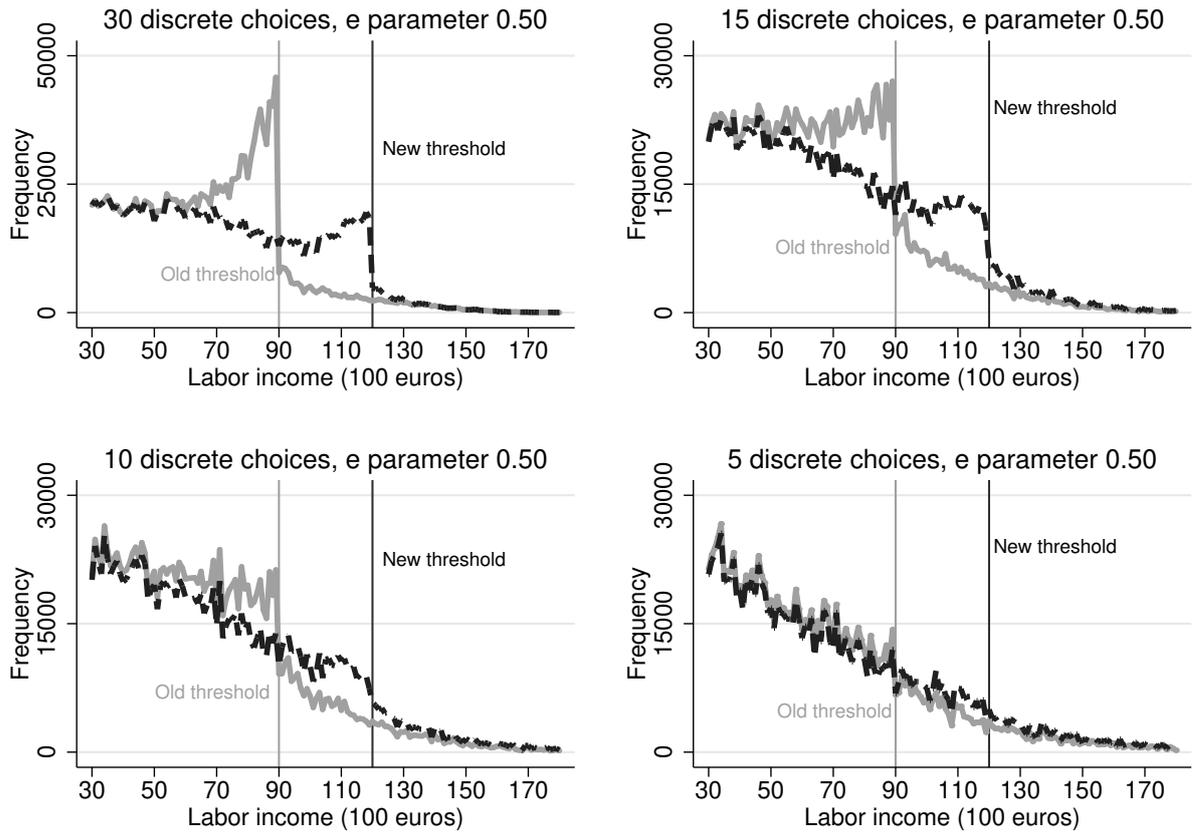
(b) Adjustment frictions



(c) Earnings shocks and adjustment frictions

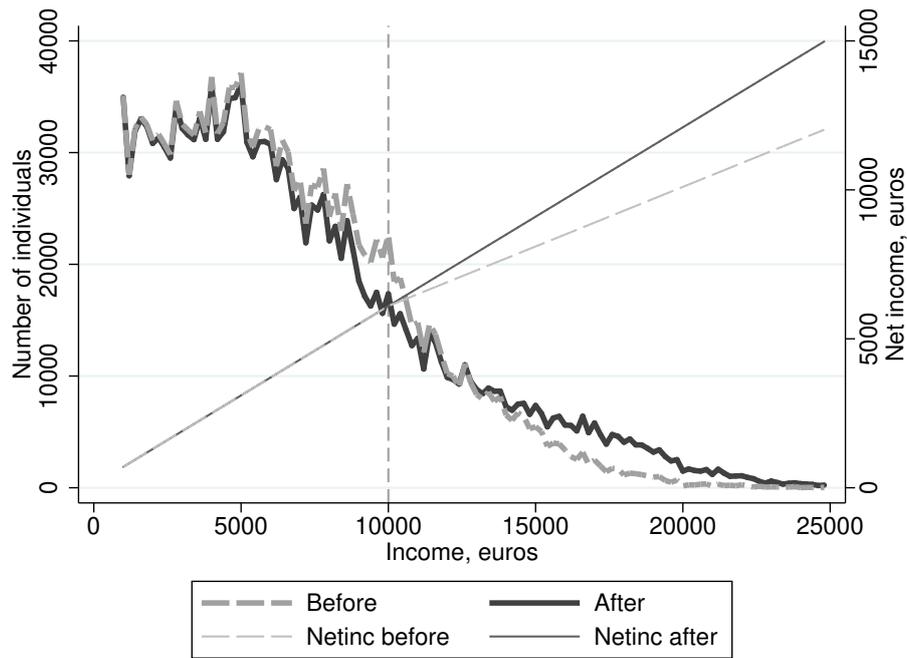
Notes: Figure presents simulated earnings distributions before (gray solid line) and after (black dashed line) an increase in the location of the notch from 9,000 euros to 12,000 euros within an income range of 0–18,000 euros. The underlying ϵ -parameter of 0.2 is used in the simulations. Qualitative results are not sensitive to the choice of this parameter value, except that with higher parameter values the densities above the thresholds reduce. Graph (a) presents the standard continuous-choice model. Graph (b) presents the standard model with adjustment frictions that prevent some students from responding to the income threshold. Graph (c) includes both adjustment frictions and unexpected i.i.d shocks in earnings to the standard model. The graphs illustrate that these frictions typically discussed in the literature can induce mitigated and scattered bunching around the threshold, but they do not produce broader changes in the earnings distributions we observed in Figure 3.

Figure 10: Simulated earnings distributions in the baseline continuous choice model and with different types of optimization frictions



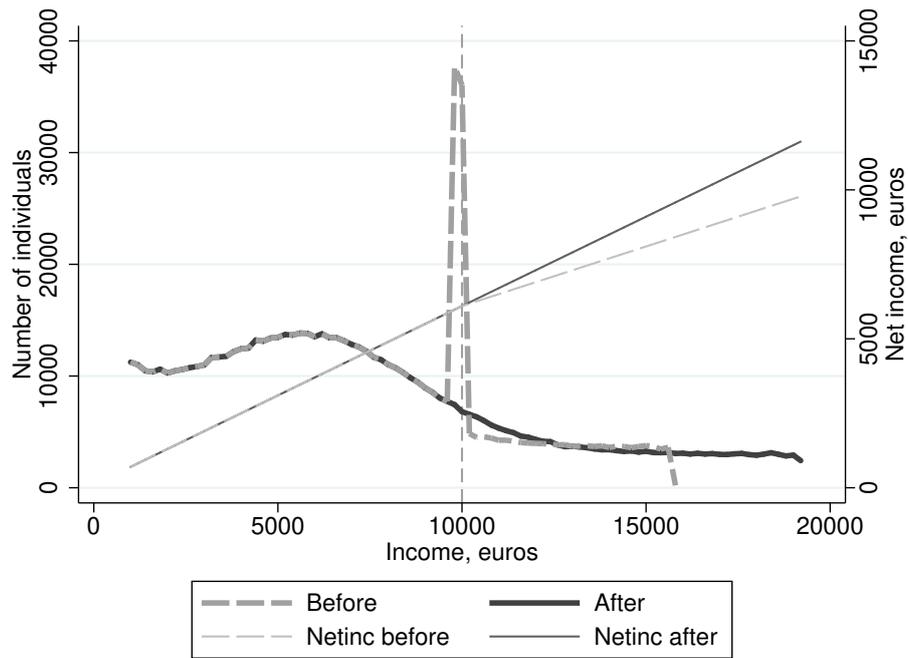
Notes: Figure presents simulated earnings distributions before (gray solid line) and after (black dashed line) an increase in the location of the notch from 9,000 euros to 12,000 euros within an income range of 0–18,000 euros using different options for an available discrete earnings choice set. The underlying e parameter of 0.5 is used in the simulations. Using 30 earnings choices produces large local bunching at the threshold, and limited changes in the distribution at lower income levels. In contrast, using 15 or 10 discrete choices produce more limited local bunching and more prevalent responses at lower income levels, similarly as in Figure 3. However, using only 5 available choices reduces both local responses and broader changes in the distribution, which is inconsistent with the empirical observations.

Figure 11: Simulated earnings distributions with different discrete earnings choice sets



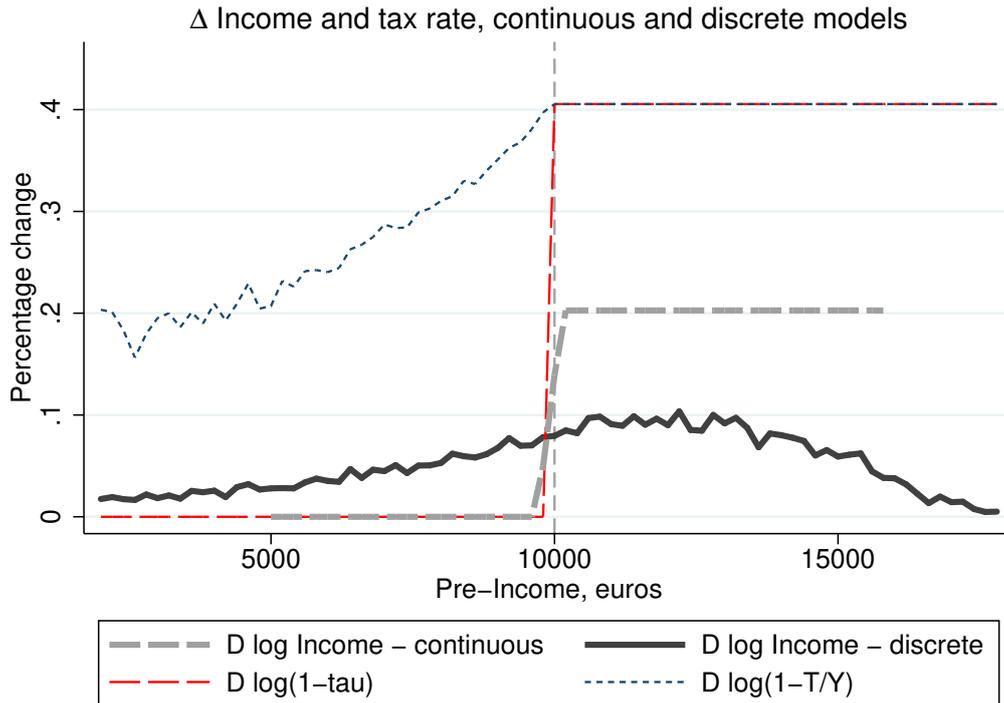
Notes: Figure presents the simulated earnings distributions before (gray dashed line) and after (black solid line) the removal of a tax rate kink at 10,000 euros (dashed vertical line). Horizontal axis denotes pre-reform income, and “Netinc before” and “Netinc after” denote the net incomes (right-hand side vertical axis) before and after the removal of the kink. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. The underlying e parameter of 0.5 and the assumption of 10 available earnings choices within 0–25,000 euros are used in the simulation.

Figure 12: Simulated income distributions with 10 discrete earnings choices before and after a removal of a tax rate kink



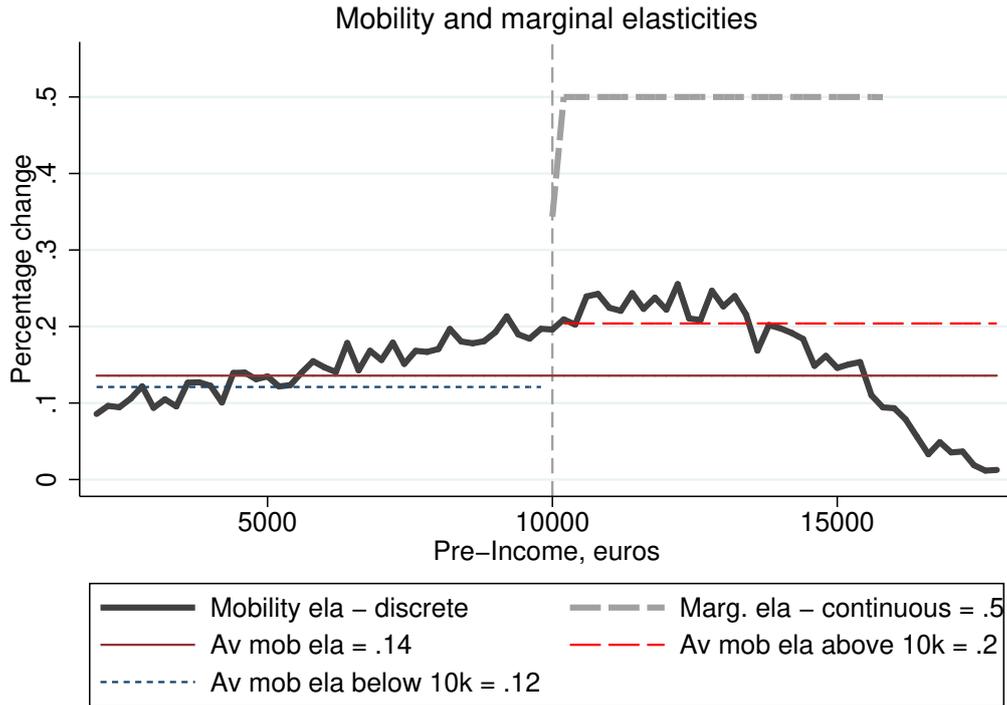
Notes: Figure presents the simulated earnings distributions before (gray dashed line) and after (black solid line) the removal of a tax rate kink at 10,000 euros. Horizontal axis denotes pre-reform income in bins of 200 euros, and “Netinc before” and “Netinc after” denote the net incomes (right-hand side vertical axis) before and after the removal of the kink. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. The underlying e parameter of 0.5 and the assumption of continuous earnings choices within 0–25,000 euros are used in the simulation.

Figure 13: Simulated income distributions with continuous earnings before and after a removal of a tax rate kink



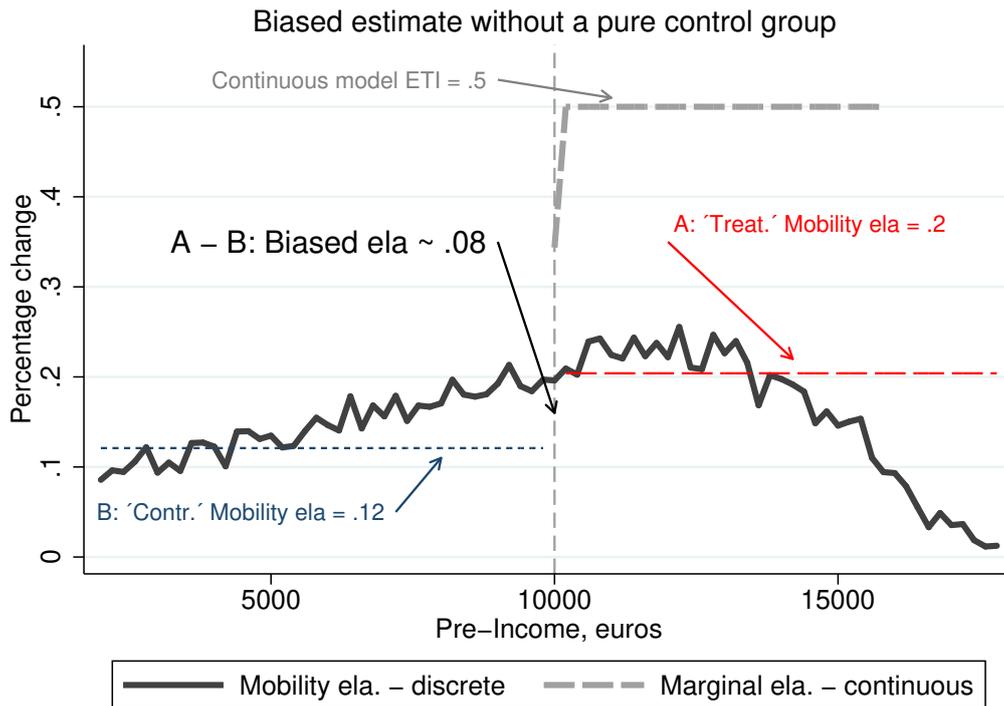
Notes: Figure presents the simulated changes in log income and tax rates for the continuous and discrete choice models after the removal of a tax rate kink at 10,000 euros using an underlying ϵ parameter of 0.5. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. Horizontal axis denotes pre-reform income in bins of 200 euros. Gray dashed line denotes the change in income in the continuous model, and solid black line in the discrete model with 10 available earnings choices within an income range of 0–25,000 euros. Red dashed line denotes the change in the net-of-tax rate in the continuous model, and the dotted blue line the change in the implied tax rate between the available earnings locations. Based on the assumptions in the continuous model, only individuals with pre-reform income above the original kink point respond to the reform, as only their net-of-tax rate is assumed to be affected in the reform. In the discrete model, a fraction of individuals in income bins below the kink also respond if their next available earnings location is above the original kink. Consequently, the implicit tax rates are also affected for those individuals below the kink, but the average change is smaller than in the continuous model above the kink. The average change in tax rate in the discrete model equals the average change in the continuous model above the kink by definition. Furthermore, in the continuous model, all individuals are assumed to respond to the reform with a similar earnings response above the kink, but in the discrete model, a smaller share of individuals respond with a discrete jump in earnings, constituting a smaller average earnings response above the kink compared to the continuous model.

Figure 14: Changes in log income and net-of-tax rates in continuous and discrete models



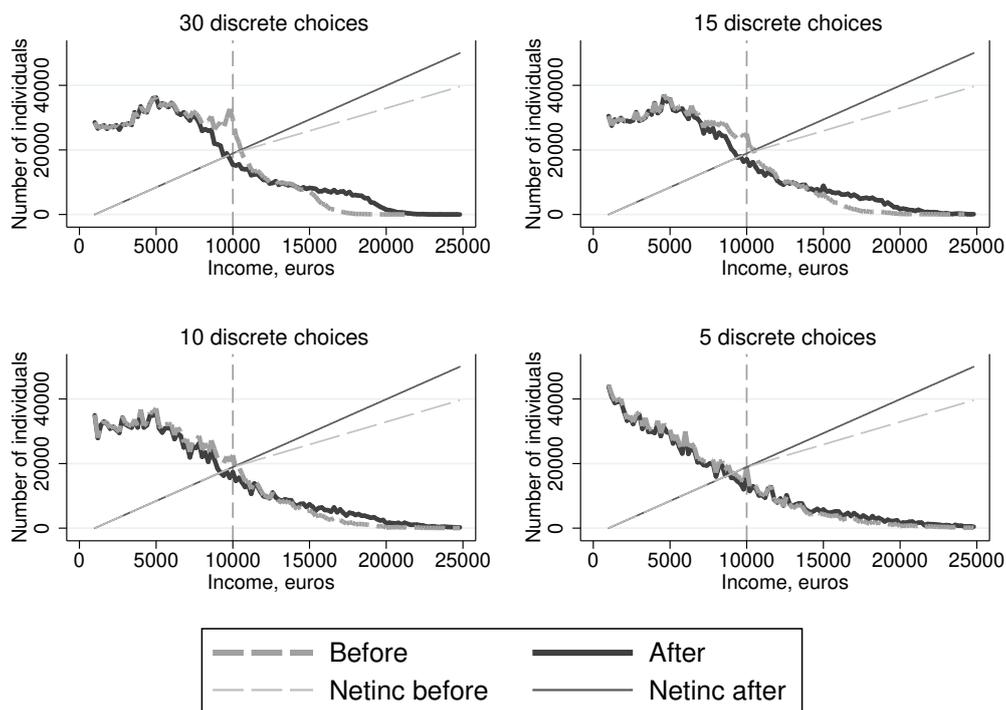
Notes: Figure presents the simulated elasticity estimates associated with the responses to the removal of a tax rate kink at 10,000 euros in discrete and continuous models using an underlying e parameter of 0.5. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. Horizontal axis denotes pre-reform income in bins of 200 euros. Solid black line denotes the average bin-level mobility elasticity estimates estimated with 10 available discrete earnings choices. Gray dashed line denotes the standard elasticity estimates from the continuous model. The standard continuous average elasticity estimate is 0.5 above the kink, which exactly equals the assumed e parameter in the model. The average mobility elasticity is 0.14, and the average estimate is 0.12 below the kink and 0.2 above it (on average). These numbers imply that when earnings choices are discrete, the true elasticity estimate can be significantly different than what we would predict based on the canonical continuous model.

Figure 15: Elasticity estimates in discrete and continuous models



Notes: Figure presents an illustration of the source of the bias in a standard differences-in-differences elasticity estimate when using those below the removed kink at 10,000 euros as a control group. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. In the canonical continuous model, the elasticity estimate is 0.5, which equals the chosen e parameter in the model. Above the original kink, the mobility elasticity estimate estimated with 10 available discrete earnings choices is 0.2 (A), and below the kink 0.12 (B). Therefore, the standard differences-in-differences approach would simply subtract the elasticity below the kink from that above it (A-B), producing a downward biased estimate of approximately 0.08. If the underlying earnings responses are continuous, no such bias emerges when using those below the kink (elasticity=0) as a control group.

Figure 16: Illustration of the bias in standard differences-in-differences estimates



Notes: Figure presents the simulated earnings distributions before (gray solid line) and after (black dashed line) the removal of a tax rate kink at 10,000 euros. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. The underlying e parameter of 0.5 and the assumption of 30, 15, 10 and 5 available earnings choices within 0–25,000 euros is used in the simulations.

Figure 17: Simulated earnings distributions with 30, 15, 10 or 5 discrete choices before and after a removal of a tax rate kink

Tables

Table 1: Descriptive statistics, all students 1999–2013

Individual characteristics					
	Age	Female	Labor income	Labor income > 500	
Mean	23.7	.56	9,130	.77	
Median	23	1	6,325	1	
sd	5.128	.496	9,524	.28	
N	5,126,594	5,126,594	4,351,213	5,126,594	
	One employer	Study subsidy months	9 subsidy months	Years studied	
Mean	.57	6.7	.32	2.1	
Median	1	8	0	2	
sd	.50	3.05	.462	1.91	
N	3,557,732	5,126,594	5,126,594	3,933,607	
Field of industry					
	Manufacturing	Services	Admin. & Publ. Sector	Other/missing	
Mean	.18	.15	.37	.29	
sd	.39	.36	.48	.45	
N	5,126,594	5,126,594	5,126,594	5,126,594	
Field of study					
	Arts & Humanities	Business & Soc. Science	Tech., Health & Soc. Serv.	Other/missing	
Mean	.13	.16	.30	.37	
sd	.33	.36	.46	.48	
N	5,126,594	5,126,594	5,126,594	5,126,594	

Notes: Table presents the descriptive statistics for all students in 1999–2013. Labor income > 500 denotes the share of students with annual labor income above 500 euros. One employer denotes the share of students who we observe to work for only one employer within a year among those with information on the employer in the data. 9 subsidy months denotes the share of students with the default study subsidy choice.

Table 2: Descriptive statistics, non-student part-time workers, 1999–2013

Individual characteristics					
	Age	Female	Labor income	Labor income > 500	Share with one employer
Mean	21	.56	8,318	.93	.62
Median	21	1	6,741	1	1
sd	1.710	.496	7,229	.25	.48
N	940,786	940,786	932,572	940,786	940,786
Field of industry					
	Industry	Services	Admin. & Public Sector	Other/missing	
Mean	.31	.22	.41	.06	
sd	.46	.41	.49	.24	
N	940,786	940,786	940,786	940,786	

Notes: Table presents the descriptive statistics for young, non-student part-time workers used in Figure 4 in the main text. The group of non-student part-time workers is selected to roughly match students' job and age characteristics. The non-student group comprise of individuals who we observe to have less than 12 working months per year in the data, and who are 19–24 years old. The age interval is chosen to match between the 25–75 percentile points of the students age distribution. Labor income > 500 denotes the share of individuals with annual labor income above 500 euros. One employer denotes the share of individuals who we observe to work for only one employer within a year among those with information on the employer in the data.

Table 3: Parameter values in the simulation model

Parameter	Value
<i>Marginal tax rate (τ)</i>	
Below the notch	0.22
Above the notch	0.61
<hr/>	
Size of the notch	500e
<hr/>	
<i>Virtual income (R)</i>	
Before	4,100e
After	3,600e
<hr/>	
<i>Location of the notch (income threshold)</i>	
Before	9,000e
After	12,000e

Notes: Table presents the parameter values used in the simulation model. The parameter values are selected to approximate the actual budget set faced by students under the study subsidy program.

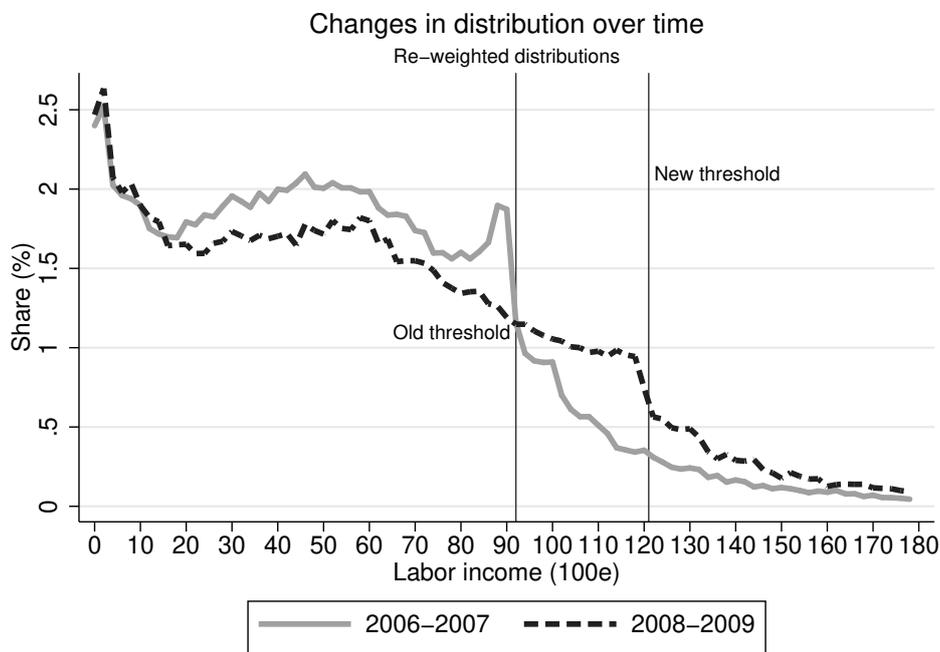
Table 4: Earnings elasticity estimates using different estimation approaches

	(1)	(2)	(3)	(4)
<i>Location of the kink</i>	kink 15k	kink 10k	kink 5k	kink 0
Naive ETI	-0.0153	0.0978	0.0833	-
Standard error	(0.000923)	(0.000839)	(0.000995)	
Mobility elasticity	0.0812	0.177	0.320	0.551
Standard error	(0.000292)	(0.000406)	(0.000563)	(0.000968)
N	3,964,692	3,959,764	3,959,070	4,000,000
R^2	0.019	0.046	0.076	0.075

Notes: Table collects the simulated earnings elasticity estimates using 10 available discrete earnings locations for each individual and an assumed ϵ parameter of 0.5, and varying the location of the kink. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink. The naive ETI estimates are estimated by regressing the change in the log of earnings on the change in the log of the net-of-tax rate using individuals below the original kink point as a control group for those individuals originally above the kink. In columns (1) and (2), individuals from a income range of 8,000 euros below the kink are used as controls, and in (3) all individuals below the kink are included. This estimate cannot be measured when the original kink is at zero pre-reform earnings in column (4). The mobility elasticity regresses the change in the log income on the change in the net-of-average tax rates between the discrete earnings locations when using a simulated group of individuals not affected by the reform as the control group. Table shows that the naive ETI estimates are downward biased compared to the underlying unbiased mobility elasticity estimates, stemming from the fact a fraction of the control group below the kink also increase their earnings in the reform. Second, the mobility elasticity estimates increase the lower is the location of the original kink in the distribution, as at lower income levels there is, on average, a larger number of available earnings locations above the kink point where individual can relocate themselves after the reform.

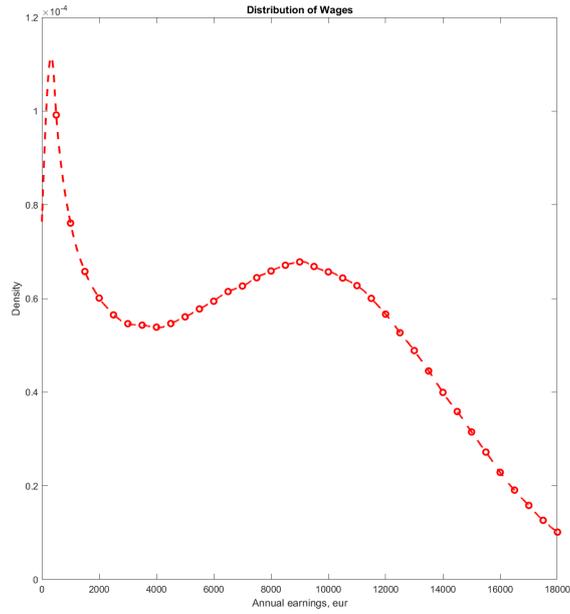
Appendix A

Figures



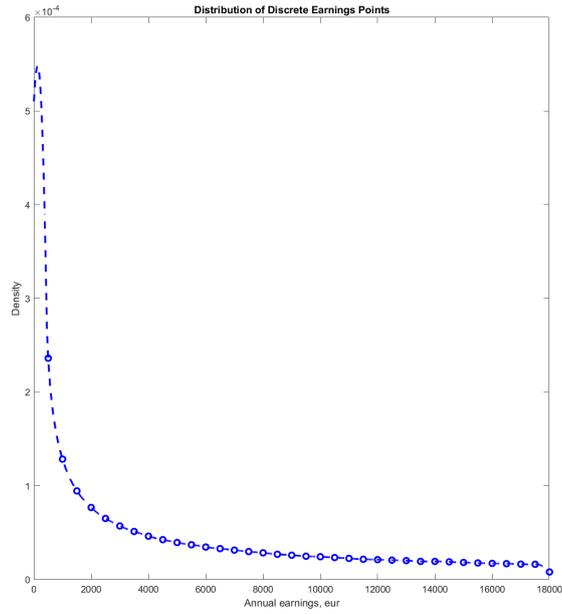
Notes: Figure presents the re-weighted observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 200 euros for students with the default 9 subsidy months in each year. Bin-level inverse probability weighting is used to re-weight the annual distributions using the year 2006 as the base year. The re-weighting procedure utilizes four groups for both the field of industry and field of study, and three age groups based on age terciles. Using equation (1), the estimate for broader changes in the distribution within an income range of 0–9,200 euros is 11.40(1.01), which is very similar to that estimated in the baseline case in Figure 3 in the main text.

Figure A1: Re-weighted earnings distributions in 2006–2007 and 2008–2009.



Notes: Figure presents the underlying earnings distribution used in the simulation model. The distribution is a combination of a power distribution and a normal distribution, which delivers an approximate match for the shape of the empirical earnings distribution of students in our empirical analysis. The simulation results are not sensitive to different underlying ability distributions that roughly match the empirical earnings distribution.

Figure A2: Simulated earnings distribution in the absence of taxes



Notes: Figure presents the underlying probability distribution of discrete earnings choices utilized in the discrete choice model simulations. The large mass in the probability distribution at small earnings ensures that each individual has at least one available choice that produces positive utility with positive earnings. The thick tail in the distribution ensures that there is another available choice at a higher income level, although the specific location of this choice can vary significantly across different draws. In the simulation procedure, we iterate the model multiple times, and in each round draw new available earnings choices. The resulting earnings distribution for the full population is continuous, although one individual faces only a discrete and limited number of possible earnings choices.

Figure A3: Probability distribution of discrete earnings choices

Tables

Table A1: Income thresholds before and after the 2008 reform

Study subsidy months	Before 2008 (academic year 2006/2007)		After 2008 (academic year 2008/2009)	
	Income threshold	Relative income loss at the margin if the threshold is exceeded	Income threshold	Relative income loss at the margin the threshold is exceeded
1	17,340	3.1%	22,550	2.5%
2	16,330	3.2%	21,190	2.7%
3	15,320	3.5%	19,930	2.9%
4	14,310	3.7%	18,620	3.1%
5	13,300	4.0%	17,310	3.3%
6	12,290	4.3%	16,000	3.6%
7	11,280	4.7%	14,690	3.9%
8	10,270	5.2%	13,380	4.3%
9	9,260	5.7%	12,070	4.8%

Note: Table presents the annual income thresholds in euros for different subsidy months before and after the 2008 reform. The highlighted 9 months of the subsidy is the default choice. The relative income loss from marginally exceeding the income threshold is calculated using the full study subsidy (461 euros and 500 euros before and after 2008, respectively) plus 15% interest collected by the Social Insurance Institution when the subsidy is reclaimed.

Table A2: Variables used in the mobility elasticity estimation for students

	Avg. gross earnings (2,000–18,000e)	Avg. net income below notch (2,000–9,300e)	Avg. net income above notch (9,300–18,000)	Differences in net incomes
2007	7,116	8,693	12,173	3,534
2008	7,529	8,785	13,592	4,807
		Gross earnings below: 6,008	Gross earnings above: 11,821	

Notes: Table presents the variables used when calculating the mobility elasticity estimate for students in Section 4.2 in the main text. Mobility elasticity is measured by relating the log change in average gross earnings to the log change in the difference of the net incomes between the two earnings locations (below and above the notch). The net earnings are calculated using the SISU-microsimulation model.

Table A3: Simulated earnings elasticity estimates using different discrete earnings choices and the original kink point at 5,000 euros

	(1)	(2)	(3)	(4)	(5)
<i># of discrete choices</i>	5	10	15	20	30
Naive ETI	-0.133	0.0833	0.197	0.263	0.337
Standard error	(0.0017)	(0.0001)	(0.0008)	(0.0006)	(0.0005)
Mobility elasticity	0.273	0.320	0.356	0.375	0.402
Standard error	(0.00096)	(0.0006)	(0.000434)	(0.0004)	(0.0003)
N	3,886,794	3,959,070	3,980,080	3,988,928	3,995,258
R^2	0.020	0.076	0.145	0.218	0.371

Notes: Table collects the simulated earnings elasticity estimates using different available discrete earnings locations for each individual (5, 10, 15, 20, 30) and an assumed e parameter of 0.5. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink located at 5,000 euros. The naive ETI estimates are estimated by regressing the change in the log of earnings on the change in the log of the net-of-tax rate using individuals below the original kink point as a control group for those individuals originally above the kink. The mobility elasticity regresses the change in the log income on the change in the net-of-average tax rates between the discrete earnings locations when using a simulated group of individuals not affected by the reform as the control group.

Table A4: Simulated earnings elasticity estimates using different discrete earnings choices and the original kink point at 10,000 euros

	(1)	(2)	(3)	(4)	(5)
<i># of discrete choices</i>	5	10	15	20	30
Naive ETI	-0.011	0.0978	0.190	0.250	0.319
Standard error	(0.0010)	(0.0008)	(0.0007)	(0.0006)	(0.0004)
Mobility elasticity	0.128	0.177	0.214	0.239	0.281
Standard error	(0.0006)	(0.0004)	(0.0003)	(0.0003)	(0.0002)
N	3,885,732	3,959,764	3,980,158	3,988,936	3,995,260
R^2	0.010	0.046	0.100	0.163	0.301

Notes: Table collects the simulated earnings elasticity estimates using different available discrete earnings locations for each individual (5, 10, 15, 20, 30) and an assumed e parameter of 0.5, and varying the location of the kink. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink located at 10,000 euros. The naive ETI estimates are estimated by regressing the change in the log of earnings on the change in the log of the net-of-tax rate using individuals below the original kink from a income range of 8,000 euros as a control group for those individuals originally above the kink. The mobility elasticity regresses the change in the log income on the change in the net-of-average tax rates between the discrete earnings locations when using a simulated group of individuals not affected by the reform as the control group.

Table A5: Simulated earnings elasticity estimates using different discrete earnings choices and the original kink point at 15,000 euros

	(1)	(2)	(3)	(4)	(5)
<i># of discrete choices</i>	5	10	15	20	30
Naive ETI	-0.0289	-0.0153	0.0194	0.0586	0.130
Standard error	(0.0009)	(0.0009)	(0.0009)	(0.0008)	(0.0006)
Mobility elasticity	0.0580	0.0812	0.0987	0.115	0.143
Standard error	(0.0005)	(0.0003)	(0.0002)	(0.0002)	(0.0002)
N	3,907,916	3,964,692	3,981,306	3,989,246	3,995,290
R^2	0.004	0.019	0.041	0.074	0.159

Notes: Table collects the simulated earnings elasticity estimates using different available discrete earnings locations for each individual (5, 10, 15, 20, 30) and an assumed ϵ parameter of 0.5, and varying the location of the kink. The marginal tax rate below the kink is 40% and 60% above it before the removal of the kink located at 10,000 euros. The naive ETI estimates are estimated by regressing the change in the log of earnings on the change in the log of the net-of-tax rate using individuals below the original kink from a income range of 8,000 euros as a control group for those individuals originally above the kink. The mobility elasticity regresses the change in the log income on the change in the net-of-average tax rates between the discrete earnings locations when using a simulated group of individuals not affected by the reform as the control group.

Appendix B

Estimating local bunching.

Behavioral responses to local discontinuous changes in the budget set, such as tax rate kinks or notches, are in the recent literature predominantly estimated using a bunching methodology (see Kleven (2016) for a summary). Intuitively, if a discontinuous jump in incentives affects earnings, we should find an excess mass of individuals located just below the threshold in the earnings distribution. This local excess bunching thus captures the total earnings distortions created by the threshold in the absence of optimization frictions. Saez (2010) and Kleven and Waseem (2013) show that under certain restrictions and within the continuous earnings supply model, the local bunching measure can be translated into an average earnings elasticity, representing a relevant parameter for the welfare analysis of taxes and income transfers.

We measure local responses to the notch caused by the income threshold following the standard bunching approach in Kleven and Waseem (2013). The local counterfactual density is estimated by fitting a flexible polynomial function to the observed distribution, excluding an area around the study subsidy income threshold z^* from the observed income distribution. We group students into income bins of 100 euros and then estimate a counterfactual density by excluding the region $[z_L, z_H]$ around the threshold from the regression:

$$c_j = \sum_{i=0}^p \beta_i (z_j)^i + \sum_{i=z_L}^{z_H} \eta_i \cdot \mathbf{1}(z_j = i) + \varepsilon_j \quad (11)$$

where c_j is the count of individuals in bin j , and z_j denotes the income level in bin j . The order of the polynomial is denoted by p . Thus the fitted values for the counterfactual density are given by $\hat{c}_j = \sum_{i=0}^p \beta_i (z_j)^i$. The local excess bunching is then estimated by relating the actual number of students close to the threshold within (z_L, z^*) to the estimated counterfactual density in the same region:

$$\hat{b}(z^*) = \frac{\sum_{i=z_L}^{z^*} (c_j - \hat{c}_j)}{\sum_{i=z_L}^{z^*} \hat{c}_j / N_j} \quad (12)$$

where N_j is the number of bins within $[z_L, z^*]$.

Following Kleven and Waseem (2013), we set the lower limit of the excluded region (z_L) based on visual observations of the income distribution to represent the point in the distribution where the bunching behavior begins, i.e. when the density begins to increase. We determine z_H such that the estimated excess mass, $\hat{b}_E(z^*) = (\sum_{i=z_L}^{z^*} c_j - \hat{c}_j)$, equals the estimated missing mass above the threshold, $\hat{b}_M(z^*) = (\sum_{i=z^*}^{z_H} \hat{c}_j - c_j)$, stemming from individuals originally above the income threshold who respond to the notch by bunching below it. We apply this convergence condition by starting from a small value of z_H and increasing it gradually until $\hat{b}_E(z^*) \approx \hat{b}_M(z^*)$. This convergence condition also defines the marginal buncher student with income $z^* + \Delta z$, representing the student with highest earnings in the absence of the notch who responds by locating below the income threshold.

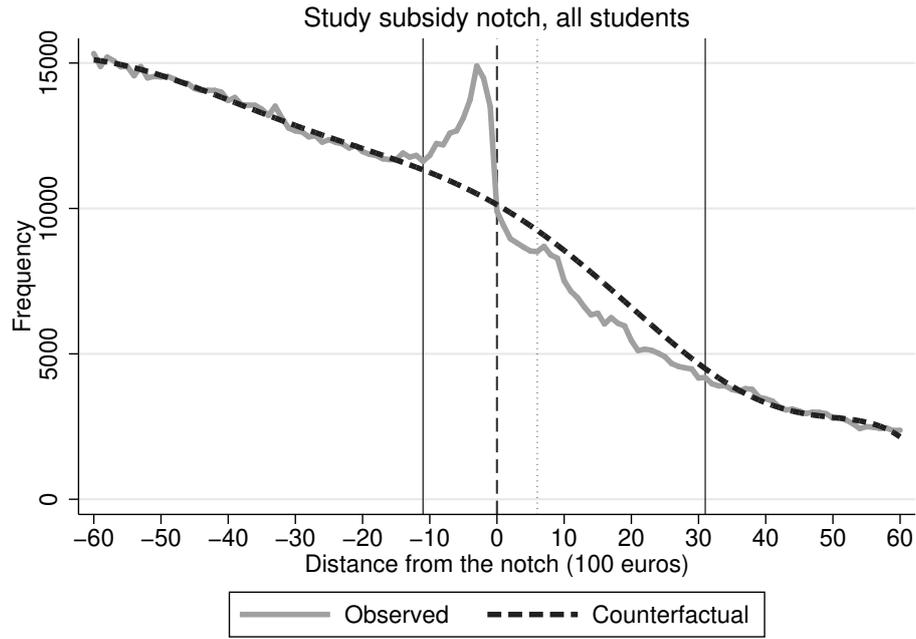
Following Kleven and Waseem (2013), we calculate standard errors by using a residual-based bootstrap procedure. We generate a large number of income distributions by randomly resampling the residuals from equation (11) with replacement, and generate a large number of new estimates of the counterfactual density based on the resampled distributions. Based on the bootstrapped counterfactual densities, we evaluate variation in the bunching estimate. The standard error is defined as the standard deviation in the distribution of the estimate.

Local bunching responses.

We find clear local responses to the income threshold of the study subsidy program. Figure B1 shows the gross income distribution and the counterfactual distribution relative to the notch in bins of 100 euros in the range of $\pm 6,000$ euros from the notch in 1999–2013. The dashed vertical line denotes the notch point above which a student loses one month of the subsidy. The solid vertical lines denote the excluded range used in the estimation of the counterfactual, which is estimated using a 7th-order polynomial function. The dash-point vertical line above the notch shows the upper limit for the dominated region just above the notch where students can increase their net income by lowering their gross income.

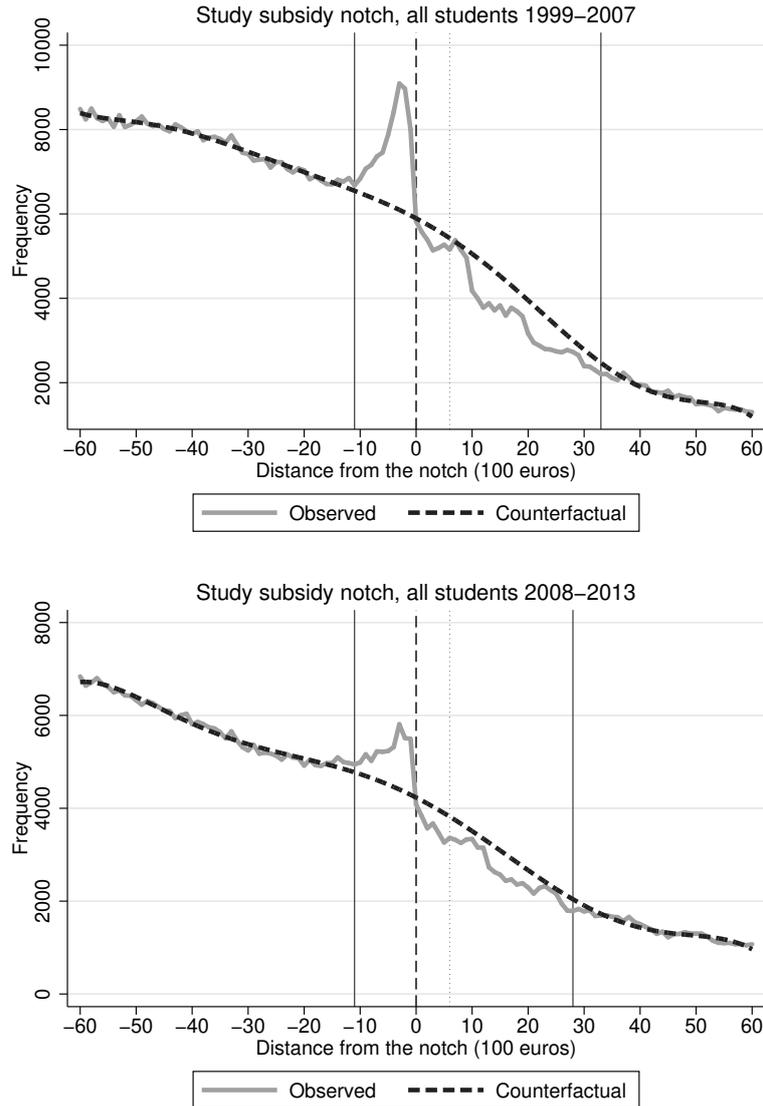
Figure B1 indicates a visually clear and statistically significant excess mass (2.19(0.189)) below the income threshold for all students. This implies that students are both aware of the notch and respond to the strong local incentives created by it. In addition, there is clear evidence of the existence of some types of frictions. There is an economically and statistically significant mass of students, 0.915(.027) of the mass relative to the counterfactual, at the locally dominated region just above the notch where no students should locate in the absence of any types of frictions (Kleven and Waseem 2013). Furthermore, even though the study subsidy schedule ultimately consists of multiple notches, we observe a distinctive response only to the first income threshold they face.

Figure B2 shows the local bunching responses before (1999–2007) and after (2008–2013) the 2008 reform. The figure shows that local excess bunching is slightly larger before (2.55(0.228)) than after (1.71(0.882)) the reform. One explanation for this is that local incentives not to exceed the notch are somewhat smaller after 2008, since the relative significance of losing one month's subsidy in terms of disposable income is now smaller than before 2008 when the threshold was at a lower income level. However, as discussed in Section 5 in the main text, the local bunching method is not a valid measure for estimating behavioral responses to tax incentives under the discrete earnings constraint, and therefore these estimates need to be interpreted as suggestive.



Notes: Figure presents the observed earnings distribution (gray solid line) and the estimated counterfactual distribution (black dashed line) around the income threshold (denoted by zero in the figure) in bins of 100 euros for all students using pooled data from 1999–2013. The first and second solid vertical lines denote the lower and upper limits of the excluded region when estimating the counterfactual distribution. The counterfactual is estimated using a seventh-order polynomial. The dotted vertical line denotes the upper limit of the region of dominated choice just above threshold. The estimate for local excess bunching at the notch is 2.19(0.189), and the estimate for the mass at the dominate region is 0.915(0.027).

Figure B1: Local bunching at the study subsidy notch, 1999–2013



Notes: Figure presents the observed earnings distributions (gray solid line) and the estimated counterfactual distributions (black dashed line) around the income threshold (denoted by zero in the figure) in bins of 100 euros for all students before (1999–2007) and after (2008–2013) the 2008 threshold reform. The first and second solid vertical lines in the figure denote the lower and upper limits of the excluded region when estimating the counterfactual distribution. The counterfactual is estimated using a seventh-order polynomial. The dotted vertical line denotes the upper limit of the region of dominated choice just above threshold. The estimate for local excess bunching at the notch before 2008 is 2.55(0.228) and 1.71(0.882) after the reform.

Figure B2: Bunching at the study subsidy notch: Before and after the 2008 reform

Local earnings elasticity estimates.

We can also estimate a local earnings elasticity estimate at the income threshold. We approximate the earnings elasticity at the study subsidy notch using a similar approach as Kleven and Waseem (2013). We derive an upper-bound reduced-form earnings elasticity by relating the earnings response of a marginal buncher student at z^H to the implicit change in tax liability between the notch point z^* and z^H (see Figure 2 in the main text). The marginal buncher represents the individual with the highest income to move to the notch point, compared to a

counterfactual state in the absence of the notch. Intuitively, this approach treats the notch as a hypothetical kink which creates a jump in the implied marginal tax rate. More formally, the reduced-form earnings elasticity is calculated with a quadratic formula

$$e(z^*) \approx (z^H/z^*)^2/(\Delta t/(1-t)) \quad (13)$$

where $(1-t)$ is the net-of-tax rate at the notch, and Δt defines the change in the implied marginal tax rate for the marginal buncher with an earnings response of Δz .

The implied earnings elasticities are 0.083(0.019) for all students and 0.065(0.007) for students with 9 subsidy months (standard errors in parenthesis). Nevertheless, as discussed above, the local bunching measure does not capture all earnings responses when earnings choices are discrete, and therefore these estimates do not represent the true earnings elasticity of students. We discuss this issue in more detail in Section 5 in the main text.¹⁹

¹⁹Furthermore, the region of dominated choice above a notch point is not necessarily a sub-optimal choice for an individual with a discrete earnings constraint. Therefore, following the approach in Kleven and Waseem (2013) and relating the share of individuals in the dominated range to the estimated local counterfactual does not necessarily give us a robust measure for other frictions affecting local responses to taxes used to approximate the structural earnings elasticity in the absence of adjustment frictions or optimization errors.