Discrete Earnings Responses: Empirical Evidence and Implications

Tuomas Kosonen and Tuomas Matikka*

August 16, 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 discrete earnings constraints provide a feasible explanation for why we typically observe modest tax elasticity estimates among wage earners in various countries and contexts.

Keywords: discrete earnings choices; tax incentives; optimization frictions
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) and CESifo, tuomas.matikka@vatt.fi.

1 Introduction

In all developed countries, a major part of the population receive their income through wage earnings, thus being subject to various income tax schedules that often take complex forms. Therefore, a question of how taxpayers' earnings respond to these tax systems has received extensive attention in the economics literature (see e.g. Saez *et al.* (2012) for a survey). In recent empirical studies, the focus has turned to understanding factors that prevent individuals from responding to taxes. Much of this work has focused on various optimization and adjustment frictions that are built on the canonical continuous earnings supply model (see e.g. Chetty *et al.* (2011), Chetty (2012), Chetty *et al.* (2013), Chetty and Saez (2013), Kleven and Waseem (2013), and Gelber *et al.* (2018)). However, labor market frictions that hinder continuous earnings adjustments to tax rate changes, such as discrete earnings choices, have received much less attention in the empirical literature, even though the idea of discrete labor supply responses is not novel (see e.g. Dickens and Lundberg (1993), Saez (2002) and Blundell *et al.* (2008)). Nevertheless, the distinction between continuous and discrete choice models is essential, as it has various key implications in explaining observed earnings responses to taxes and their welfare considerations. For example, the sufficient statistic for the welfare analysis of taxes is the observed elasticity of taxable income (ETI) estimate under discrete earnings, while under continuous model and adjustment frictions the observed estimate deviates from the structural elasticity.

In this paper, we provide new and transparent quasi-experimental evidence that give clear support for discrete earnings responses among wage earners, and analyze the broader implications of this fundamental labor market friction. We utilize a novel empirical design to uncover the mechanisms behind observed earnings adjustment: a combination of an income notch creating strong local tax incentives, and a reform that shifted out the location of the notch. Using the relocation of the notch, we can detect responses to a local incentive from a broader income range, revealing us how the notch affects the earnings choices far away from the discontinuity. Furthermore, we develop theoretical arguments and provide illustrative simulation results highlighting the key role of discrete earnings choices as the driving mechanism. Furthermore, we show that discrete choices instead of continuous earnings adjustments provide a feasible explanation for why the ETI estimate is typically observed to be rather low among wage earners.

Despite the extensive literature on earnings responses to taxes, there is only limited evidence on the mechanisms behind the earnings adjustment process, presumably due to the lack of suitable quasi-experimental set-ups to credibly study this question. In order to uncover these mechanisms, we study the earnings behavior of Finnish university students receiving a monthly study subsidy. A 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 a clear income notch for students above which their disposable income reduces sharply. In 2008, the income threshold was increased by 30%, allowing students to earn more income before they lose the subsidy they are eligible for.

The income threshold and its relocation provide an excellent set-up to study the mechanisms behind the earnings responses. First, the notch creates a large local change in incentives, potentially creating large earnings responses that are more feasible to detect in empirical data than responses to smaller incentives. Second, the reform that shifts out the location of the notch allows us to observe potential earnings responses far away from the threshold that we otherwise could not identify from the earnings distribution in a stationary notch context.

Moreover, university students are an excellent population to study the presence of frictions stemming from the labor market. In Finland, university students typically participate in flexible part-time and temporary labor markets during their studies within and between semesters. They often work for many employers within a year, creating further flexibility in responding to incentives. Therefore, any discreteness affecting the earnings responses of students is also likely to be present for other taxpayers in a more permanent labor market context. 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.

To estimate the effects of the reform that shifted out the location of the notch, we develop novel methods that calculate the change in the relative density in the distribution associated with the reform. The estimation operates in the same spirit as the methods that calculate the excess mass in the bunching method, but in this case calculate the change in the density in a wider range in the income distribution between two time periods. To be able to take into account general changes in the labor markets affecting income of everyone, we also subtract from this estimate the change in the income mass in a control group not affected by the study subsidy system. In this case the control group consists of young people in the part-time labor markets resembling university students in these characteristics.

We find that the change in the location of the notch caused clear earnings responses in a broad income range. After the reform, a visible excess mass immediately below the old notch disappeared and a new, somewhat smaller, excess mass appeared just below the new location of the notch. More importantly, we observe that many students who were located far below the old notch significantly increased their earnings precisely at the time of the reform. 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 incomes. To support the notion of causality, we do not detect any changes in the shape of the earnings distribution of our control group, other young part-time workers at the time of the reform. This indicates that the earnings responses of students did not arise from other contemporary changes in the part-time labor market in Finland. Furthermore, the location of the earnings distribution of students remains practically constant in other years than 2008, further supporting the causal impact of the income threshold 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, at the time of the reform in 2008, students originally below the income threshold were much more likely to increase their earnings by 50% or more, compared to the period before 2008. These results support the notion that the earnings responses within a broad income range arise because of individual-level relocation responses rather than, for example, extensive margin responses through a more productive students participating in the labor market after the reform.

More generally, earnings responses within a broad income range well below the notch indicate that students responded to the reform by discrete earnings responses rather than continuous earnings adjustments suggested by the canonical labor supply model. Under the traditional continuous choice model, we would have expected just the local excess mass to be relocated from below the old threshold to below the new threshold, which indeed has also occurred. However, we would have not expected other significant changes in the shape of the distribution that we nonetheless observe. This line of thought is further supported by the fact that we find clearly larger shifts in the income distribution of students working in more inflexible and discrete labor markets (any public sector, or private sector industries: research, manufacturing and construction), compared to students working in labor markets that typically have more flexibility and less discreteness in labor supply choices (restaurants, bars and cafes, hotels, cleaning and security services).

We develop a simple theoretical framework and provide illustrative simulation results to examine the underlying mechanisms in more detail. We find that a standard labor supply model is not able to explain the empirical results. In our simulations, this model generates extensive and sharp bunching just below the income thresholds, but no earnings responses from below the notch when the threshold is relocated. Furthermore, adding adjustment frictions or optimization errors discussed in the more recent literature do not change this main intuition. Intuitively, in the continuous choice model, an individual originally well below the old notch would not change her earnings as a response to an increase in the location of the notch to a higher income level. The original location choice was determined by optimizing her earnings with respect to the tax rate schedule and preferences, and the inputs in this optimization process for this individual did not change in the reform.

However, when the available earnings choices are limited to a discrete choice set, tax rate changes far away from the original earnings location can affect the choices of an individual, that is, can induce large earnings responses. Therefore, the discrete choice model can rationalize much larger jumps in earnings as a response to local tax rate changes compared to traditional labor supply models. Indeed, adding the discrete earnings choice constraint in the simulation model generates earnings distributions that largely share the key characteristics with their empirical counterparts. When using a relatively sparse choice set of 10–20 available choices (on average, jumps of 1,250–2,500 euros in annual earnings), we can qualitatively match two key empirical features: shifting of the income distribution from a wide income range and scattered local bunching.

We surmise that discrete earnings choices for individuals arise because their earnings arise from the labor markets and there it is not necessarily possible to make marginal earnings adjustments. There are many reasons for this. Empirically, detailed Finnish register data on wage rates and working hours show that underlying individual earnings decisions often include discrete components such as fixed hourly or monthly wages coupled with fixed or restricted monthly working time. These are stemming from, for example, working hours regulations and collective agreements between the labor market parties. More broadly the institutions in the labor markets are such that employees and employers need to make contracts on the specifics of job description, and these contracts have a fixed length or one or several months notice periods, and thus cannot be changed continuously.

This study contributes to the extensive literature examining earnings and labor supply responses to tax incentives. In particular, we show that discrete earnings constraints can provide a viable explanation for why the estimated ETI tend to be rather low among regular wage earners, and why the estimates tend to vary between different countries and contexts (see e.g. Saez *et al.* (2011) for a survey). We use our simulation framework to illustrate that discrete earnings constraints have a significant impact on the observed differences-in-differences estimates of the tax rate elasticity. For example, when assuming an underlying earnings disutility parameter such that ETI estimate is 0.7 in the continuous case, but limiting the available individual earnings choices to 10 (average jumps of 20–25% in earnings), the estimated elasticity reduces to 0.21. Also, when implementing the discrete earnings constraint, the R^2 statistics is reduced to the level we typically observe ($<$0.1) when empirically estimating the causal impact of income taxes on earnings. Therefore, the discrete earnings constraint can provide a reasonable and straightforward explanation for why we typically observe modest earnings elasticities among wage earners in many contexts.

In addition, we contribute to the literature discussing the welfare implications of taxes (see e.g. Chetty (2009, 2012)). As the discrete earnings constraint in pure form are persistent for individuals, the observed elasticity estimates under this optimization friction represent a relevant parameter for welfare analysis. Earlier literature has focused on optimization frictions that attenuate responses to tax incentives, such as inattention, salience of tax rules and regulations, optimization errors and other types of earnings adjustment costs (Chetty *et al.* 2009; Chetty *et al.* 2011; Chetty *et al.* 2013; Chetty and Saez 2013; Kleven and Waseem 2013; Gelber *et al.* 2018; Gelber *et al.* 2019). A recent and related paper in this field is Søgaard (2019), who studies the bunching responses of university students in Denmark to a kink created by the Danish study subsidy system. He finds that inattention to incentives appears to be the dominant friction. In the presence of these adjustment-type optimization frictions the observed elasticity estimate could be substantially lower than the structural elasticity estimate that is the sufficient statistic for welfare analysis. Thus, the distinction between our paper and the earlier studies is that the types of optimization frictions analyzed have clearly different welfare implications. Our paper and these earlier papers show together, that empirically both types of frictions, adjustment and discrete, are likely to be relevant.

Furthermore, we contribute to the bunching literature estimating 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. For example, Blomquist *et al.* (2018) argue that observed bunching at kinks and notches does not provide evidence of the elasticity of taxable income without explicit assumptions on the distribution of heterogeneity. Our contribution is that the discrete earnings choices limit the applicability of local bunching estimates differently from the other reasons discussed in recent literature. First, under discrete earnings choices, the bunching method underestimates the extent of behavioral responses. Second, discrete earnings responses far away from the discontinuity affect the shape of the distribution, and consequently, the surrounding distribution outside the bunching region cannot necessarily be used to estimate a credible counterfactual describing the shape of the distribution in the absence of a notch or a kink. To illustrate this, we find that the local excess mass of students just below the income threshold underestimates the overall earnings responses to this notch by a factor of 3.6.

This study contributes also to several other branches of literature. Most of the previous literature discussing discrete earnings choices 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 of wage earners are discrete (see van Soest (1995) and Beffy *et al.* (2019), and Löffler *et al.* (2018) for a recent review). However, there is only limited reduced-form evidence of discrete earnings responses, perhaps due to a lack of suitable quasi-experimental set-ups 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 quasi-experimental 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 mechanisms and present our simulation models and discuss their implications. Section 5 discusses the broader implications of discrete earnings choices. Section 6 concludes the study.

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 the predetermined threshold, the study subsidy of one month is reclaimed by the SII. This results in an increase in average tax rate, or 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

[1] The 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.

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 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 is 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

[2] In 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.

[3] After 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.

received labor earnings from only one employer, suggesting that students tend to work in different types of jobs during the year.

Overall, these observation 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.

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 slightly younger (22.4) and have less labor income (5,633) than all students (on average). 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 the 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 new methods to estimate these changes that build on the 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

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

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 denotes the overall number of students and N_j denotes the number of bins within $[z_L, z_H]$.

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

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 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 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.[5] However, our differences-in-differences estimator does not require that the two groups have exactly the same pre-reform income distribution. Similarly to the standard differences-in-differences, the requirement is that the changes in the two distributions should develop similarly over time.

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

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

where superscript S denotes students and P non-student part-time workers. This estimate thus summarizes the broader change in the earnings caused by the reform while taking into account other potential changes in the labor market environment of part-time workers.[6]

[5]The 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.

[6]Following the bunching literature, the standard errors for $\hat{b_d}(z)$ are calculated using a residual-based bootstrap procedure. 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

3 Main results

In this Section we estimate the impact of the notch and the reform shifting it's location on the earnings choices of students utilizing the methods described in the previous Section.

Figure 3 shows the labor earnings density 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.[7] Remarkably, the figure shows that the earnings distribution of students has a significantly different shape in 2008–2009 than before the reform, 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, we estimate the differences-in-differences equation (2) within an income range of 0–9,200 euros, thus including the whole distribution below the 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.[8] 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 below the new threshold within the year of the reform in 2008. 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

replacement). The standard error is defined as the standard deviation of $\hat{b}_d(z)$ based on the bootstrapped distributions.

[7]The figure includes only labor earnings as receiving capital income is very rare among university students.

[8]The local bunching method and local bunching results are discussed in more detail in Appendix B.

distinctive local response only to the first income threshold.[9]

As a 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 pre-reform characteristics in terms of age, field of study and field of industry. This bin-level inverse probability weighting procedure accounts for potential differences in key student characteristics between the periods. However, re-weighting does not change the outcomes in a significant manner, indicating that potential changes in the characteristics of the student population over time are not likely to explain the results either.[10]

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 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

[9] Additional 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 smaller than before 2008 when the threshold was at a lower income level.

[10] In 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.

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 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.[11]

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 Conceptual framework

Next, we discuss different theoretical models that could or could not explain the main empirical results presented above. The main feature of the results that 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. We start with a standard continuous earnings supply model and extend that with optimization frictions and discrete choice sets.

A standard continuous framework features a utility function over consumption and

[11] Table A3 in Appendix A further summarizes the income responses of students using income transition matrices that further describe the location choices of students in the earnings distribution before and after the reform. We find that at the time of the reform in 2007–2008, the share of students who increased their earnings for more than 3,000 euros increased in income bins below the old threshold when compared to the time period before the reform, i.e. 2006–2007.

leisure and a linearized budget set consisting of earnings, consumption and income taxes. From an individual point of view, the exogenous preferences and taxes determine the earnings location individual chooses in the income distribution. Formally, the utility function is $u(c, z)$, where c denotes consumption and z 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:

$$u(c, z) = c - \frac{w^i}{1 + \frac{1}{e}} \left(\frac{z}{w^i}\right)^{1+\frac{1}{e}}, \tag{3}$$

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 and describing the relative magnitude of welfare losses in an optimal income tax setting.[12] Intuitively, the income location choices of an individual i are determined by innate productivity w^i, and the response to taxes determined by the utility function. In this 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 if tax rates are changed, for example, at a higher earnings level.

An increasingly popular extension to this canonical framework is optimization frictions that attenuate behavioral responses to tax incentives (see e.g. Chetty *et al.* (2011), Chetty (2012), Chetty *et al.* (2013), Chetty and Saez (2013), and Kleven and Waseem (2013)). The optimization frictions typically considered include job switching costs, salience of tax rules or unawareness of tax incentives. The above 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 frictionless model. The utility function then becomes as follows

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

From the above equation it becomes clear that considering these kind of optimization frictions merely reduce the responsiveness to taxes, but they do not alter individual

[12] Note that parameter e captures the elasticity with respect to taxes, defined as $\frac{dz}{d(1-\tau)} \frac{1-\tau}{z} = e$, only within this parameterized utility function.

earnings responses to taxes in a more fundamental manner.[13]

We can further alter the basic framework by adding optimization errors to the model, arising from an unanticipated shock to the initially chosen income. The simplest 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 choice is $z^* - r$. A crucial aspect of optimization errors is that they are unanticipated and thus do not enter the decision model determining the optimal income choice. Therefore, these kind 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.

Next, we consider a model that can explain broader changes in earnings as a response to a change in local tax incentives; discrete earnings choices. We define discrete earnings choices as having a non-continuous and relatively small number of alternative earnings locations from which the individual must choose from. One motivation for a discrete earnings choice set is, for example, that individuals can typically choose between a limited number of different employers who offer jobs with a fixed monthly salary, constituting a rather sparse choice set. In addition, earnings adjustments within jobs can include discrete wage rate and working hours opportunities, leading to an non-continuous earnings choice sets. Furthermore, the idea of discrete earnings choices is not new in the conceptual literature. For example, they were part of an explanation offered for the lack of bunching at tax rate kinks in Saez (1999), Lundberg and Dickens (1993) analyzed them theoretically, and structural labor supply models often assume that labor supply choices (working hours) are discrete in nature (see e.g. van Soest (1995), and Löffler *et al.* (2018) for a recent review).[14]

Following Saez (2002), discrete earnings could be modeled through a constraint 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. In the model, individuals must now 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 wage distribution are continuous as before. Thus, the utility function would be as above, but indexing the discrete earnings and consumption choices with j, denoting the available earnings choices for an individual i. The discrete choice is determined by which location, $j-1$ or j yields higher utility.

$$u(c_{j-1}, z_{j-1}) \leq u(c_j, z_j) = c_j - \frac{w^i}{1+\frac{1}{e}} \left(\frac{z_j}{w^i}\right)^{1+\frac{1}{e}} \tag{5}$$

[13]Heterogeneous adjustment costs do not change this general intuition, but would explain why simultaneously some taxpayers respond to taxes while others do not.

[14]We discuss the empirical literature on discrete earnings choices below in Section 4.3.

The conceptual difference between this model and the canonical 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 schedule occurs far away from the current location in the earnings distribution. In other words, tax rate changes far away from the current location can induce earnings responses. The discrete choice model can therefore rationalize much larger jumps in earnings as a response to a local tax rate change, compared to any of the continuous models considered above.

Figure 6 illustrates the differences between continuous and discrete choice models when a tax notch changes its location. 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. 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. Moreover, the indifference curves do not need to tangent the budget set in the discrete choice model. Additionally, in the discrete model, the region above the notch is no longer necessarily strictly dominated, as depending on the discrete earnings choices that are available, it is possible that the best available discrete point is located just above the notch.

To summarize, our analysis above 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.

4.2 Simulation results

Baseline model and optimization frictions.

Next, we utilize the conceptual framework above and provide simulation results which further illustrate the income responses under different model assumptions. We begin by illustrating how students would respond to the relocation of the income threshold in the baseline continuous choice model and in the presence of typical optimization frictions and optimization errors. In order to have a budget set that resembles the empirical one for students under the study subsidy system, we assume parameters given in Table 3. The marginal income tax rate is set to 22% below the notch. To simplify the analysis, we assume a high marginal tax rate of 61% above the notch which approximately linearizes the budget set with many subsequent notches above the income threshold, presented in Figure 1 above. The size of the notch, i.e. the size of the drop in disposable income at the threshold, is approximately 500 euros. The model utilizes an underlying ability distribution that translates into an income distribution in the absence of taxes or other constraints. Each individual receives a predetermined draw from this distribution. Our parameterized ability distribution is presented in Figure A2 in Appendix A.[15]

Graph (a) of Figure 7 shows the simulated distributions when using the baseline continuous model (equation (3)) and a baseline earnings elasticity parameter of 0.2.[16] The two distributions in the figure correspond to those simulated using the budget set before (solid gray line) and after (dashed black line) the relocation of the income threshold from 9,200 to 12,000 euros. Bunching at the income threshold is sharp and sizable in both cases, in contrast to more attenuated and diffuse bunching in the empirical distribution (Figure 3). Most importantly, the canonical continuous model does not produce any changes in the earnings distribution further away from the notch point, as already discussed above in Section 4.1.

Next, we add optimization frictions to the model following equation (4). We assume heterogeneous frictions represented by a uniformly distributed parameter a in the unit interval. Each individual has a different and independent draw from this distribution. Graph (b) of Figure 7 shows that adding these frictions alters the shape of the distribution only by reducing the size of the spikes at the notch, and inducing some individuals to be located in the dominated range above the threshold. Furthermore, we add unexpected earnings shocks by including i.i.d. normal distributed mean-zero income shocks with a standard deviation of 800 euros to the model in graph (c) of Figure 7. The figure illustrates that local bunching is now more diffuse, as in the empirical distribution.[17] However, these simulated distributions still do not capture the broader changes in the overall earnings distribution, implying that adding optimization frictions or errors is not sufficient to explain our empirical findings.

Discrete earnings choices.

Next, we add discrete earnings choices as an additional constraint to individuals' behavior. We utilize the baseline framework presented above but limit the earnings choices of each

[15]The 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.

[16]Qualitative implications of this model are not sensitive to the elasticity parameter, except that with larger values the density above the thresholds reduces.

[17]If we were to assume only negative income shocks, we would get diffuse bunching only below the notch, similarly as in the empirical distribution. However, modeling shocks such that they affect earnings only to a certain direction is more difficult to justify from a theoretical perspective.

individual to a discrete set of available earnings choices. In the model, the discrete earnings choices are drawn from a power distribution weighting lower income options but spanning over the whole income range, presented in Figure A3 in Appendix A.[18] We iterate the model multiple times, and in each round we draw new available earnings choices. Thus, the resulting earnings distribution for the full population is continuous, although one individual faces only a limited number of possible earnings choices.

Figure 8 illustrates the earnings distributions using 30, 15, 10 and 5 available choices for each individual within an income interval of 0–25,000 euros and assuming an underlying earnings elasticity parameter of 0.7. With a discrete choice constraint included in the model, the earnings distributions and the response to the 2008 reform begin to resemble the empirical distributions of students (Figure 3). In particular, when the number of discrete choices is set to 15 or 10 (upper-left and lower-right graphs), 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. Note that scattered local bunching below the threshold only results from including the discrete earnings constraint, i.e. this model does not include optimization errors or other frictions. Overall, these findings clearly support the key role of discrete earnings choices in explaining our empirical findings.

To further support the notion that the assumption of discrete earnings choices governs the observed changes in the earnings distribution, Figure 9 shows the simulated distributions when fixing the number of available discrete choices to 10 but varying the assumption on the underlying elasticity parameter from 0.4 to 1.25. The figure shows that the overall shape of the distribution and broader income responses after the relocation of the notch are rather similar irrespective of the assumption of the elasticity parameter. However, assuming an elasticity parameter towards the higher end of this range increases the broader income responses below the notch and reduces the mass fairly significantly in the upper tail of the distribution, which are not in line with the empirical distributions.

Table 4 collects the estimates for broader changes in the simulated earnings distribution when using different assumptions on the number of available discrete earnings choices and the underlying elasticity parameter in the model. In the estimations, we set the upper limit below the old location of the notch at 7,700 euros because we primarily want to capture changes in the distribution below the local bunching region. Thus, we estimate equation (1) in the income range of 0–7,700 euros. We then compare these estimates to those we observe in the empirical distribution within the same income interval. In the table, we highlight the simulations with 10–20 available discrete earnings choices and an

[18]The large mass in the probability distribution at small earnings ensures that each individual has at least one choice that gives 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 across different draws.

underlying elasticity parameter within a range of 0.6–0.8, as these simulations produce similar estimates for the shifting of the distributions as their empirical counterpart.

These results further underline that we need a clearly discrete choice set of in order to match the key characteristics of the empirical distribution. This implies that, on average, the discrete earnings choice set compatible with the empirical findings includes 10–20 available earnings choices that are, on average, 1,250–2,500 euros apart from each other. However, none of model specifications discussed above provides a complete match with the empirical distributions. Nevertheless, our primary goal was not to build a model that would completely replicate the empirical observations, but rather to describe which elements would be crucial for any model capable of explaining the mechanisms behind our main empirical results.

4.3 Further support for discrete earnings choices

The results above illustrate that discrete earnings constraints are a significant factor even for Finnish university students who typically participate in the flexible part-time and temporary labor markets. Next, we further motivate and present empirical support for the discrete earnings model in a more general context. In general, even though we typically observe more or less continuous aggregate earnings distributions in many countries and contexts, the underlying individual earnings decisions often include discrete components. This claim is based on the fact that a major part of annual earnings stem from hourly or monthly wages that are multiplied by the annual time spent working. Both of these elements, wage rates and working hours, typically include discrete or discontinuous elements at the individual level.

This discreteness arises from at least two institutional factors. First, in many types of labor markets, wage rates and working hours are regulated either directly by legislation or by collective agreements between the representatives of employees' and employers' organizations. For example, in most developed countries, there are minimum wage regulations and working hours restrictions in place in various occupations, which translate into discreteness in individual-level earnings choices. Furthermore, collective agreements between the labor market parties set the wage rates and working hours for various duties in many occupations for a fixed time period, implying that wage rate and working hours choices within an occupation can include significant inflexibility and discreteness.

Second, the discrete nature of wage rates and hours also arise from the employment contracts between workers and firms. Even in the absence of restrictive legislation or collective agreements, both employees and employers can benefit from fixed longer-term contracts that set either the wage rate or working hours for a given period of time, or both. From the workers point of view, fixed contracts ensure a predictable level of future earnings. From the employer's side, they ensure a sufficient labor force for a fixed time period and reduce bargaining costs. These contracts often imply that wage earners commit to, for example, a full-time job for a set time period, or the available choices for part-time employment include only a limited number of available hours opportunities. This feature of many labor markets further underline the discrete nature of available earnings choices for wage earners. Furthermore, as these types of constraints stemming from regulation and employment contracts are more or less permanent in nature, the discrete earnings constraints are likely to play at least some role in labor supply responses also when considering longer-run responses to taxes and other incentives.

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 detailed occupation-level data on wage rates and working hours. First, Figure 10 presents two pieces of descriptive evidence of the discreteness in wage rates at certain industries. 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 collective agreements between the labor parties. Therefore, from the individual point of view, wage rate increases (or reductions) often occur in a discontinuous manner. Similarly, the cleaners wage rate distribution relative to the personal minimum wage presented in graph (b) illustrates that most workers in that field receive the minimum wage or a wage rate very close to it, and thus the overall wage distribution is characterized by one discrete spike, and a small amount of observations at other wage rates. Therefore, minimum wage regulations can induce significant constraints in terms of wage rate responses for many workers.

Second, individual working hours follow a discontinuous pattern in many countries and contexts. Graph (a) in Figure 11 shows the weekly working hours for all wage earners in Finland, highlighting that in many cases employment contracts commit workers for a full-time job for a set time period, which can be seen as a very distinctive spike in full-time working hours (typically 37.5 or 36 hours per week in Finland). In graph (b), we exclude full-time work from the figure to underline that there are also clear focal points in the distribution of part-time work stemming from the work-time regulation and typical part-time employment contracts, such as 18 or 30 hours per week.

In addition to a more general empirical motivation for discrete earnings choices, we next illustrate how students working in different types of jobs responded to the 2008 income threshold reform. This heterogeneity analysis provides additional suggestive evidence on the role and relevance of discrete earnings constraints in different types of labor markets. In Figure 12, we present the earnings distributions before and after the 2008 reform for students working in labor markets with presumably less discrete available choices, including restaurants, bars and cafes, hotels and other accommodation services, cleaning and security services, and retail sales such as supermarkets and gas stations.

As a suggestive comparison group, we present the distributions for students working in labor markets with a presumably more discrete available choice set, including the public sector and research, and manufacturing and construction. Employers in the less discrete job group typically offer hourly-paid vacancies and more flexibility in affecting individual working hours, compared to, for example, typical public sector or manufacturing jobs. Therefore, we would expect that students working in the less discrete group have more opportunities to adjust their labor supply and earnings compared to students working in the other group, and we would then expect to observe less discrete earnings responses to the relocation of the notch for the less discrete labor market group.

Indeed, Figure 12 provides evidence that the broader changes in the earnings distribution after 2008 are smaller for those students who work in less discrete labor markets (6.14(1.71)), and more prevalent for those working in jobs with less available discrete choices (10.94(1.10)). However, broader earnings changes due to the reform are not insignificant for those in the first group either, illustrating that discrete earnings choices can induce relevant constraints even in these types of labor markets. These findings also 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 issue in more detail in Section 5.

Finally, as discussed above, the idea of contract and regulation-driven discreteness and inflexibility in the labor market is not new in the economics literature. In particular, 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. Manning (2003) discusses the role of labor market power of employers in affecting labor supply responses and working hours choices of employees. Structural labor supply models used to analyze labor supply responses and choices in various applications 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 tend to cluster at certain focal points in the distribution. Nevertheless, there is only limited causal evidence of the underlying mechanisms behind earnings and hours responses to policy changes, perhaps due to a lack of suitable quasi-experimental research set-ups and detailed individual-level data. 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.

5 Broader implications of discrete earnings choices

5.1 Welfare considerations

Next, we discuss the broader implications of discrete earnings choices in terms of both conceptual and empirical analysis of labor supply and earnings responses. We begin by briefly discussing whether discrete earnings constraints affect the welfare analysis of income taxes and social transfers, which is at the core of public finance research, aiming at relating the observed earnings responses caused by taxes to a single statistic characterizing the welfare impact of these policies. In recent literature, a key method in measuring the welfare consequences of various policies is the sufficient statistics approach (see e.g. Chetty (2009)). In this approach, 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 estimate 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 using the continuous choice model presented above, the welfare loss from a tax rate change can be formalized as follows:

$$\text{W(t)}=\left\{(1-\tau)TI-\sum_{i=1}^{N}\frac{w^i}{1+\frac{1}{e}}\left(\frac{z}{w^i}\right)^{1+\frac{1}{e}}\right\}+\tau * TI \tag{6}$$

where TI refers to total taxable income (earnings), and the individual behavior is aggregated over N individuals. At the core of the sufficient statistics approach is the idea that individuals have already optimized their behavior, i.e. they have chosen the optimal level of earnings with a given tax rate. Therefore, we can use the envelope condition and ignore the direct change in utility stemming from the individuals' earnings choices, and only focus on changes stemming from the government revenue constraint, i.e. how much taxable earnings respond to changes in tax rates. Therefore, the change in welfare is captured by the following condition:

$$\frac{dW(t)}{d\tau}=-TI+TI+\tau\frac{dTI}{d\tau} \tag{7}$$

The first two terms in the right-hand side cancel each other out, and thus the sufficient statistics for welfare analysis is the elasticity of taxable earnings.

As discussed already in Chetty (2009), replacing the continuous choice model with discrete earnings choices does not change the fundamental idea of the sufficient statistic formula. As long as individuals make optimal earnings choices given the discreteness constraints, which are assumed to be beyond the control of individuals, we can still utilize the envelope condition and arrive at a similar conclusion as above. Intuitively, assuming discrete rather than continuous earnings choices does not change the impact on the government revenue constraint, as it is defined similarly in both cases (aggregate earnings over all individuals times the tax rate).

However, the persistence of the discrete earnings constraint might matter for welfare considerations. The baseline objective of the welfare analysis is to describe the welfare losses stemming from the underlying structural earnings elasticity (see Chetty (2012) and Kleven and Waseem (2013)). Thus, we would prefer to apply an elasticity estimate that is unaffected by any short-run constraints or frictions in equation (7). Therefore, the most relevant aspect of the welfare analysis is to analyze how the available earnings choice set affects the estimated earnings elasticity. In the static discrete earnings model the discrete constraint is permanent and thus would not change over time. Empirically, we think that discreteness arises from the functioning of the labor markets, and thus at least some part of it is always present as a constraint for individual choices. This consideration is in contrast to the earlier optimization frictions literature focusing on adjustment costs, which could well be removed over time. Thus the discrete earnings model captures a very different welfare consideration than the adjustment cost models.

5.2 How would discrete earnings choices affect earnings elasticity estimates?

Earnings responses to income taxes have been extensively studied in the economics literature. Recent estimates for the elasticity of taxable income using individual-level data typically range between 0–0.5 (see Saez *et al.* (2012) for a review), more recent results often being on the lower end of this scale. In more recent empirical work, the focus has turned to understanding and uncovering factors that attenuate responses to taxes, such as salience of tax rules, inattention and earnings adjustment frictions, which we discussed and analyzed already above. However, to our knowledge, there is no comprehensive analysis on how the discrete earnings choice constraint would affect estimated elasticities. Given that discrete earnings choices are likely to be a prominent feature in many labor markets, this analysis can provide relevant and important insight on why we typically observe relatively small earnings elasticity estimates.

In general the responses to taxes in the discrete earnings model are captured by mobility elasticity (Saez (2002)). Contrary to the continuous model, where affected individuals respond marginally in their income, the mobility elasticity measures the fraction of individuals moving to their next available location as a response to change in the average tax rates. Following Saez (2002), we can define the mobility elasticity with the following equation.

$$\zeta_i = \frac{dh_i}{dT} \frac{z_i - z_{i-1}}{h_i}$$

, where dT is the marginal change in the taxes, $z_i - z_{i-1}$ the difference in after-tax incomes in the two states, h_i the density of individuals type i and dh_i the change in the density. To calculate the elasticity in our case, we could measure the change in density relative to the whole population in the region below the notch $\frac{dh_i}{h_i}$ and then change in the (inverse of) average tax rates $\frac{z_i - z_{i-1}}{dT}$. However, this requires assuming all individuals below the notch are of the same type and our tax rate with the changing notch design is not exactly marginal.

Instead, we use the simulation framework presented above to study the role of discrete earnings choices in affecting the magnitude of observed elasticity estimates. In the baseline simulations, we use a similar framework as in Section 4.2 and set up a simple hypothetical tax reform with a marginal tax rate increasing from 0 to 30% for income of 10,000 euros onward. We model individual behavior within an income range of 0–25,000 euros, and thus this reform would correspond to a large marginal tax increase for higher incomes. First, we simulate earnings data with different available earnings choice sets and a given underlying elasticity parameter, and then estimate the ETI utilizing this hypothetical tax reform and a differences-in-differences approach, where the tax rate is increased only for the treatment group in the post-reform period while there is no increase in taxes for the control group. The estimation thus utilizes simulated earnings distributions in the before and after periods for the treatment and control groups, using the difference of log earnings as the dependent variable and the difference of the log net-of-tax rate as the independent variable. Furthermore, the differences-in-differences estimate excludes an income range of 400 euros from around the kink point to avoid sharp bunching affecting the estimate in the baseline continuous choice model.

Figure 13 illustrates the changes in the simulated income distribution with 8 discrete earnings choices for each individual. Similarly as with the income threshold simulations above, the income distribution changes its shape from a broader income range around the kink in the after period, but because of the discrete choice constraint, there is no sharp bunching exactly at the kink.

Table 5 collects our baseline elasticities with different available earnings choices. In column (1), we assume an underlying elasticity parameter of 0.7 when simulating the income distributions, and use 10,000 discrete choices to approximate the continuous choice case. The estimated ETI is 0.7, which exactly corresponds to the selected elasticity parameter in the underlying model. Column (2) presents otherwise the same model but reducing the available discrete choices to 10. Consequently, the ETI estimate from the difference-in-differences estimation reduces to 0.21. In column (3), we further reduce the number of available discrete choices to 5, which produces an estimated ETI of 0.04. Column (4) presents a model where the underlying elasticity parameter in the utility function is 0.4 and the number of available discrete choices is 10. The estimated ETI in this model is 0.12, which is again well below the assumed underlying elasticity parameter used in the earnings simulations. Furthermore, another interesting observation from the table is that the R^2 statistic reduces significantly when the number of available discrete earnings choices are reduced. This reduction in the predictive power of the regression model is impressive given that we did not add any other form of heterogeneity or frictions in the model than the randomly allocated discrete earnings choices.

Our baseline simulation results offer two key insights. First, we find that the discrete earnings constraint has a distinctive impact on the estimated earnings elasticity. Even when assuming a large underlying elasticity parameter in the model, discrete earnings choices can significantly reduce the average responsiveness to tax rate changes. Second, the R^2 statistics is reduced to the level we typically observe when empirically estimating the causal impact of taxes on earnings. Therefore, the discrete earnings constraint can provide a reasonable and straightforward explanation for why we typically observe modest earnings elasticities among wage earners in many contexts, particularly in the short run.

Table 5 presents additional simulation results using a larger number of available earnings choices (sub-table (a)), and a smaller tax reform that increases the tax rate from 0 to 10% instead of 30% as in the baseline model (sub-table (b)). First, the results illustrate that the underlying elasticity parameter in the utility function and the estimated elasticity significantly differ from each other only when the number of discrete earnings choices is small enough. For example, when using 100 available earnings choices (250 euro jumps in earnings on average), the ETI estimate (0.65) is rather close to the underlying elasticity of 0.7 used in the model. However, when reducing the number of discrete choices to 30, the estimate reduces to 0.51, and further reducing the number to 10 delivers an estimate of 0.21. These findings could thus offer an explanation for why we tend to observe different earnings elasticities in different contexts. If the labor market we study has more constraints in terms of earnings choices, we are likely to observe much smaller elasticity estimates. In particular, this is more likely to be the case in contexts where there are more limited opportunities to affect taxable earnings by other means than changes in labor supply or effort. In other words, pure reporting responses to taxes, such as tax avoidance and evasion, are typically associated with less constraints in earnings choices. This could at least tentatively explain why these types of responses are more commonly linked to larger observed earnings elasticities.

Second, we find that reducing the size of tax rate variation produces slightly smaller but economically similar ETI estimates. In other words, the size of the variation in incentives does not significantly affect the estimates under the discrete choice constraint. Chetty (2012) shows that under various types of adjustment frictions that attenuate responses to taxes (inattention, salience of tax rules, search costs etc.), larger variation in tax incentives is more likely to reveal an estimate for the underlying long-run structural elasticity. Intuitively, when changes in tax rates are large enough, more individuals respond to taxes as their gain from responding exceeds their adjustment cost. However, the size of the variation in incentives does not have a similar effect on the discrete earnings constraint as it does not alter the available choice set. Nevertheless, larger tax rate variation does induce more individuals to change their earnings locations, but according to our simulation results, this does not necessarily have a significant impact on the magnitude of the observed earnings elasticity.

To summarize, we find that discrete earnings constraints can have a significant impact on estimated earnings elasticities, particularly when the available earnings choices are sparse. As many labor markets are associated with earnings constraints that can be relatively permanent in nature, this constraint could at least explain the relatively small observed elasticity estimates in many previous studies, and consequently, smaller implied welfare consequences of taxes and social transfers.

5.3 Implications for local bunching estimation

The bunching model and its implications rely on the continuous choice model (Saez 2010; Kleven and Waseem 2013). However, the discrete earnings model sets some limitations to the bunching method. Given that our empirical analysis has demonstrated that earnings responses might indeed face discrete constraints due to the functioning of labor markets, thus potentially affecting all wage earners, this limitation might be important. The limitations of the local bunching approach can be illustrated using Figure 6 above, which shows that under the discrete choice constraint the relocation choices of individuals can occur on a far greater income range, especially far below the local discontinuity, than they would occur under the continuous model. Therefore, local bunching responses to a discontinuity in the budget set have a limited potential to capture all relevant responses to tax incentives when the earnings choices are discrete.

Furthermore, this implies that the surrounding distribution outside the bunching region cannot necessarily be used to estimate a credible counterfactual describing the shape of the distribution in the absence of a notch or a kink. 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 the notch or kink is relocated, we have no straightforward manner to evaluate how the surrounding density is affected by this discontinuity. 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 relevant parameters for welfare analysis in context of labor market outcomes.

In addition, the region of dominated choice just above a notch point is not necessarily a sub-optimal choice for an individual with a discrete earnings constraint. Intuitively, 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.

Finally, the elasticity estimate in the bunching method relies on the continuous labor supply model. When the labor supply model is discrete, this model cannot be used to calculate the relationship between change in tax rates and the excess mass and thus derive the elasticity. This applies also to models that add an adjustment friction parameter to the standard labor supply model to capture the adjustment frictions. Discrete earnings model yield more fundamentally different predictions, and thus the elasticity estimate using a wrong model would have limited information value.

More broadly, discrete earnings choice sets could offer us a viable explanation for why many papers do not find that regular wage earners would bunch at kink points in the tax schedule (see e.g. Saez (2010) and Bastani and Selin (2014)). In the cases where bunching responses to kink points are observable, the estimates for wage earners tend to be very small (see e.g. Chetty *et al.* (2011)). Discrete earnings choices can explain these results, as it is feasible that no discrete choices are available near the kink point resulting in scattered responses in wide income range. We illustrate this with our simulation model in Figure 13, where individuals respond to a kink point when earnings choices are discrete. The figure does not show almost any bunching at the kink, even though the kink introduces a large tax rate change from 0 to 30%. In fact, without the simulated distribution from the before period representing a clean comparison distribution in the absence of the kink point, local bunching might be difficult to detect in the distribution from the after period.

Furthermore, 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 found to be very small or zero (Saez 2010), but the differences-in-differences estimates for ETI are typically significantly larger even for regular wage earners (see e.g. Saez *et al.* (2011) and Weber (2011)). Similar evidence is also available for Sweden (see Bastani and Selin (2014) and Blomquist and Selin (2010)).

Even though the local bunching approach has some limitations when earnings choices are discrete, it can be a useful method in other contexts. For example, bunching estimates can deliver relevant evidence of the behavioral responses among entrepreneurs or in detecting tax avoidance or evasion, because in these instances discrete earnings is not a likely constraint. Furthermore, heterogeneity in local bunching responses can provide valuable insight on mechanisms explaining taxpayers' behavior.

5.4 Relating the excess mass and broader income choices of students

We are interested in to what extent quantitatively the local excess mass approach could underestimate the broader income responses of students. One approach is to relate the broader changes in the income distribution of students to the local bunching estimate derived utilizing the standard bunching method. We estimate and show in Figure 3 that the broader changes in the earnings distribution are approximately 3.3 times larger than the local bunching estimate (9.81 vs. 2.93).

An alternative approach to relate the local and broader changes is to estimate the average income lost due to the notch using only the local response or all responses. We find that the average reduction in earnings is approximately 510 euros when utilizing broader changes in the income distribution before and after 2008. This number is calculated as the difference between real average bin-level earnings of students in 2006–2007 and in 2008–2009 within an income range of 9,200–18,000 euros, thus outlining the average income loss caused by the old threshold. For the local approach, we find that the notch reduces earnings by approximately 140 euros. This number is calculated as the difference between the bin-level average earnings for the observed distribution utilizing as the counterfactual distribution the post-reform distribution. The calculation is done in the income interval between the notch (z^*) and the upper limit of the counterfactual density (z_H), which also defines the response of the marginal buncher – the last individual who responds by relocating below the threshold because of the notch in the budget set (see graph (a) of Figure 2 and Appendix B for more details). Intuitively, this number thus describes how many euros students are earning less on average because they choose to reduce their earnings from above the threshold to just below the notch point. When comparing these two measures, we can approximate that limiting the analysis to local responses would underestimate the overall earnings response by a factor of 3.6.

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 in the form of discrete 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 responses.

Discrete earnings choices have important broader implications. First, if the ETI estimates are lower because of discrete earnings choices as an optimization friction as opposed to adjustment costs as optimization frictions, the welfare implications differ. In the former case the observed ETI estimate is the structural one whereas in the latter case the observed ETI estimate could be lower than the structural ETI, which the sufficient statistic for welfare losses. Second, our simulation results highlight that discrete earnings adjustments instead of continuous earnings adjustments can provide a feasible explanation for why we tend to observe modest tax elasticity estimates among wage earners in various countries and contexts. We find that adding a discrete earnings constraint produces significantly smaller observed elasticity estimates even when the underlying disutility from work is relatively large. Third, the local bunching method produces downwards-biased earnings response estimates when earnings choices are discrete, implying that that tax elasticity estimates derived from local changes in tax rates do not necessarily provide us with applicable tax elasticity estimates. Our analysis of Finnish university students show that the local bunching responses underestimate overall earnings responses by a factor of 3.6.

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] Kleven, H. (2016). Bunching. Annual Review of Economics, Vol. 8, 2016.

[17] 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.

[18] Löffler, M., A. Peichl and S. Siegloch (2018). The Sensitivity of Structural Labor Supply Estimations to Modeling Assumptions. ifo Working Papers 259/2018.

[19] Manning, A. (2003). Monopsony in Motion: Imperfect Competition in Labor Markets. Princeton, NJ: Princeton University Press.

[20] Saez, E. (1999). Do Taxpayers Bunch at Kink Points? NBER Working Papers 7366, National Bureau of Economic Research.

[21] Saez, E. (2002). Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses, Quarterly Journal of Economics, 117: 1039–1073.

[22] Saez, E. (2010). Do taxpayers bunch at kink points? American Economic Journal: Economic Policy, 2(3): 180–212.

[23] 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.

[24] Søgaard, J. (2019). Labor supply and optimization frictions: Evidence from the Danish student labor market. Journal of Public Economics, Vol. 173: 125–138.

[25] van Soest, A. (1995). Structural models of family labor supply - a discrete choice approach. Journal of Human Resources, 30(1): 63–88.

[26] 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

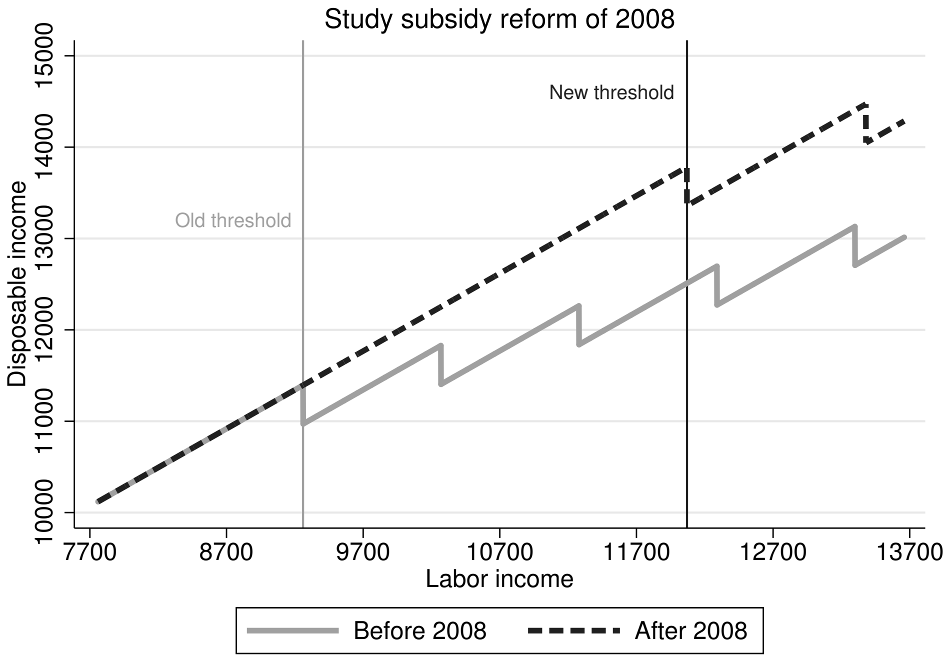

Notes: Figure presents the study subsidy schedule before (gray solid line) and after 2008 (black dashed line) for a student who collects the default 9 subsidy months. The vertical axis denotes disposable income, and horizontal axis labor income. The vertical lines denote the thresholds before (9,200 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

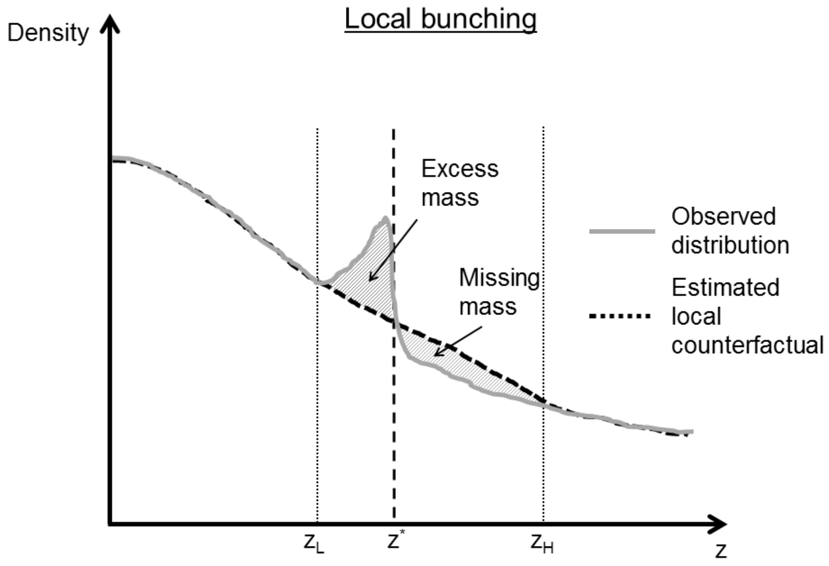

(a) Local bunching

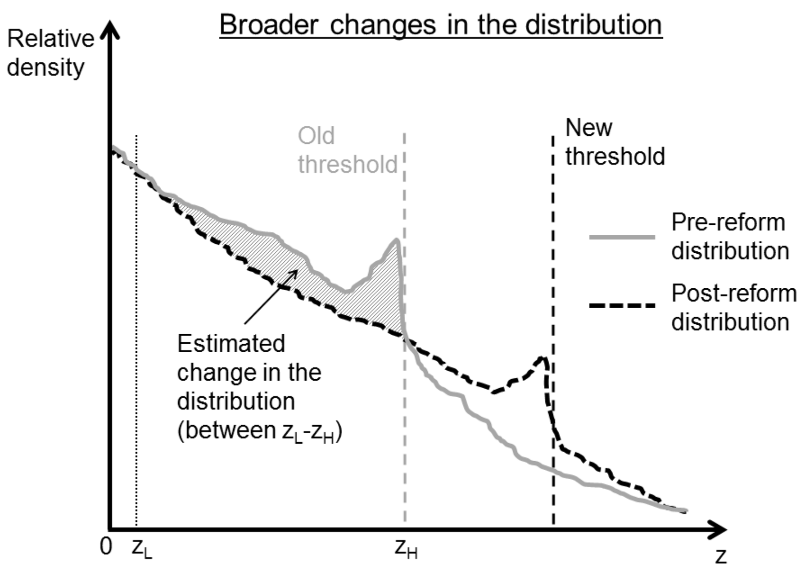

(b) Broader changes in the distribution

Notes: Graph (a) illustrates the local excess bunching at the income threshold in a hypothetical earnings 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.

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

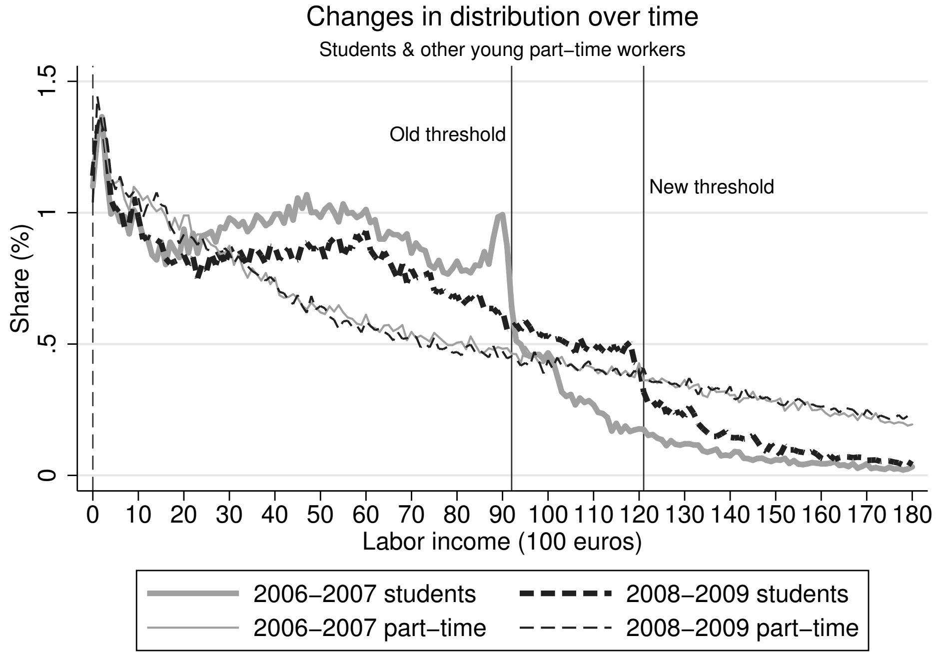

Notes: Figure presents the observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 100 euros for students with the default 9 subsidy months in each year and for 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

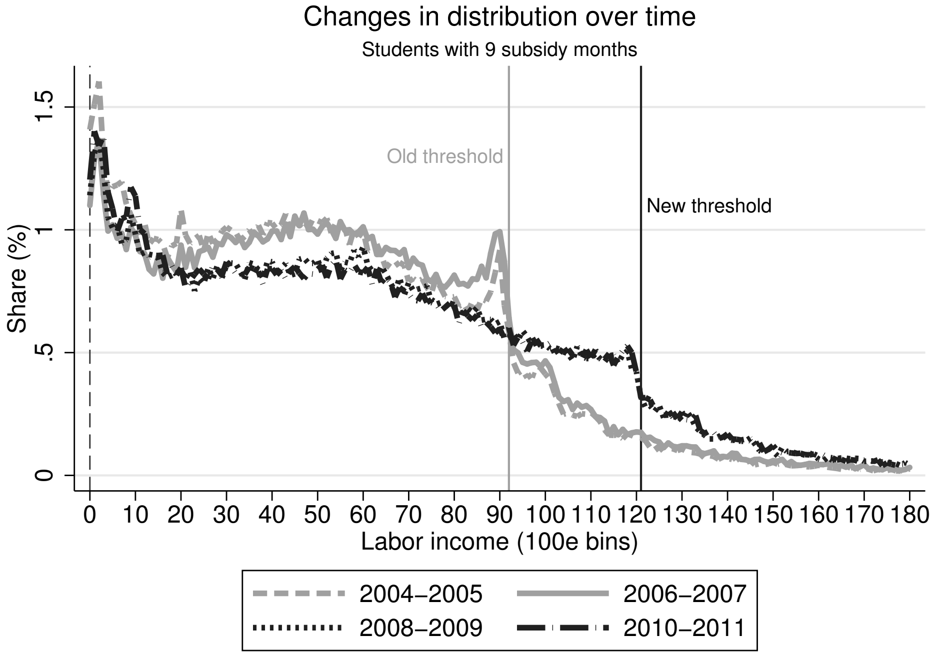

Notes: Graph (a) presents the earnings distributions of students with 9 subsidy months in 2004–2005 (gray dashed line), 2006–2007 (gray solid line), 2008–2009 (black solid line) and 2010–2011 (black dotted line) within an income range of 0–18,000 euros in bins of 100 euros. The figure shows that the response of students occurred exactly at the time of the 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

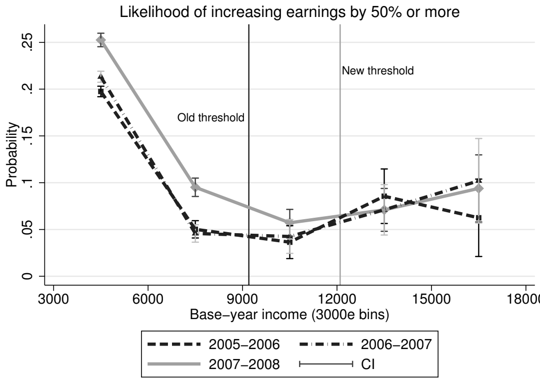

(a) More than 50% earnings increases

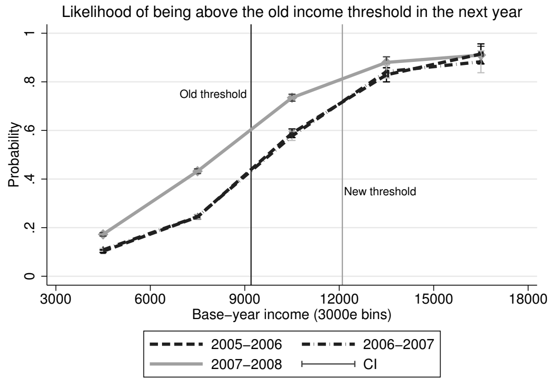

(b) Locating above the old income threshold

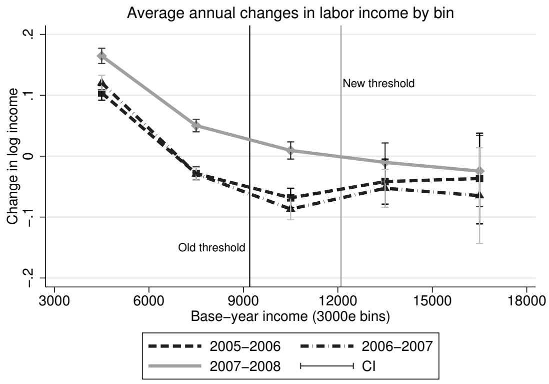

(c) Average changes in labor income

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 are significantly larger 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 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

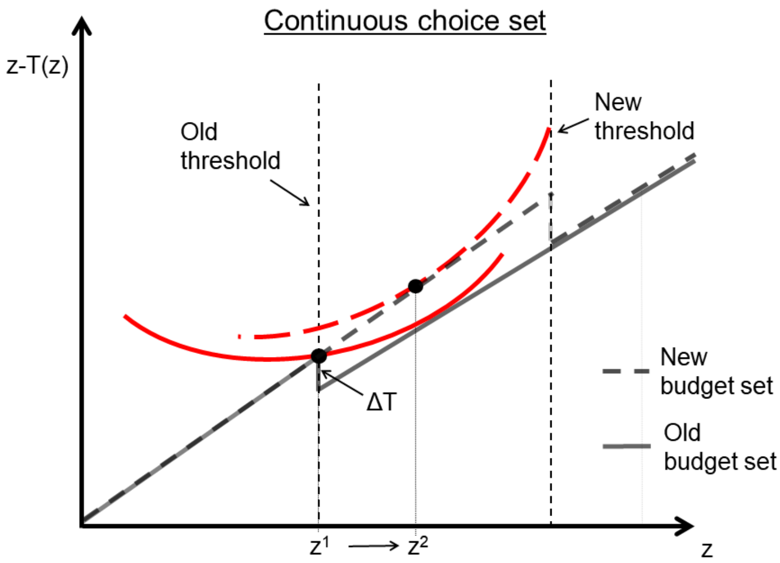

(a) Continuous choice model

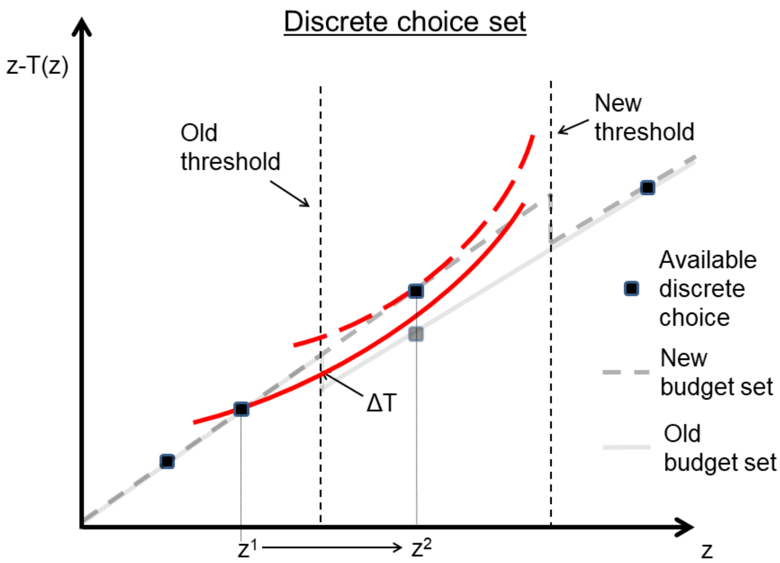

(b) Discrete choice model

Notes: Figure illustrates the theoretical earnings responses to a relocation of a tax notch in hypothetical budget set diagram. Graph (a) shows the case of fully continuous earnings choices. Graph (b) illustrates the case where an individual can only choose from a limited set of available discrete earnings choices. The graphs highlight that under the discrete choice constraint, we are more likely to observe earnings responses from below the notch to above it.

Figure 6: Relocation of the notch in a continuous earnings model and a in a discrete choice model

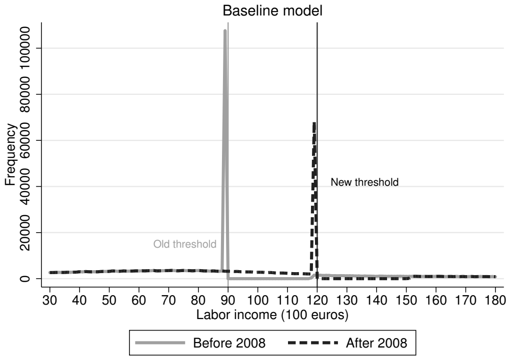

(a) Baseline simulation model

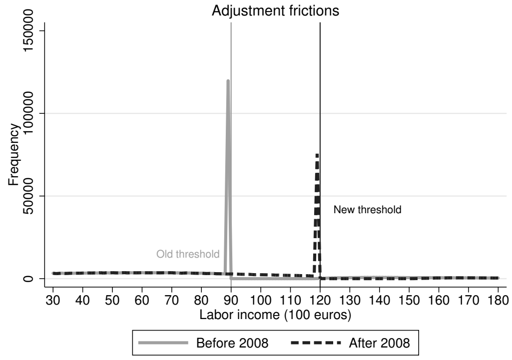

(b) Adjustment frictions

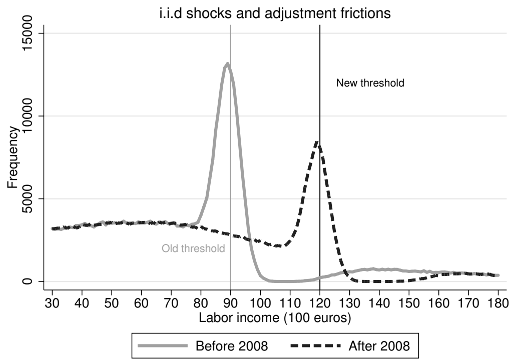

(c) Earnings shocks and adjustment frictions

Notes: Figure presents simulated earnings distributions before (gray solid line) and after (black dashed line) the 2008 reform within an income range of 0–18,000 euros. The underlying elasticity parameter of 0.2 is used in the simulations. Qualitative results are not sensitive to the choice of this parameter value, except that with larger values of the elasticity parameter the densities above the thresholds reduce. Graph (a) presents the standard continuous choice model with no optimization frictions. 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 that are 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 7: Simulated earnings distributions in the baseline continuous choice model and with different types of optimization frictions

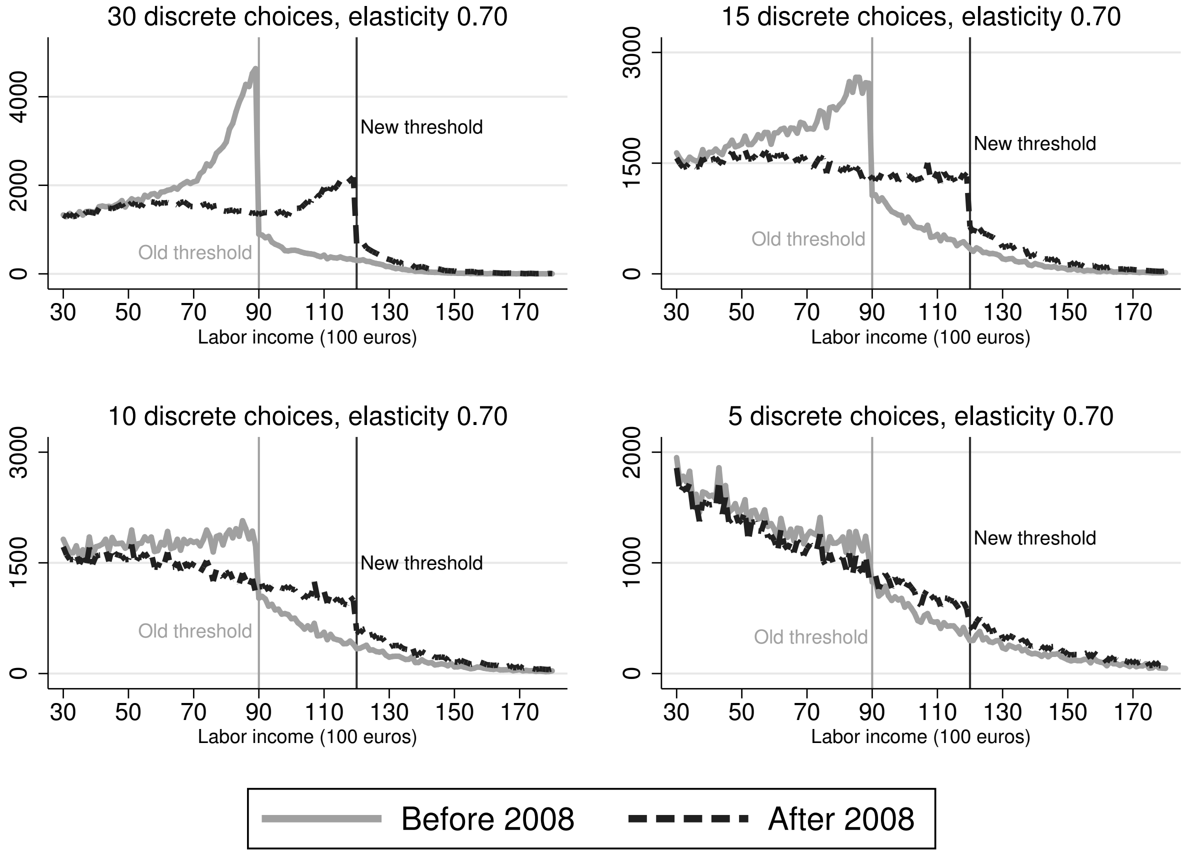

Notes: Figure presents simulated earnings distributions before (gray solid line) and after (black dashed line) the 2008 reform within an income range of 0–18,000 euros using different options for the available discrete earnings choice set. The underlying elasticity parameter of 0.7 is used in the simulations. 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 8: Simulated earnings distributions with different discrete earnings choice sets

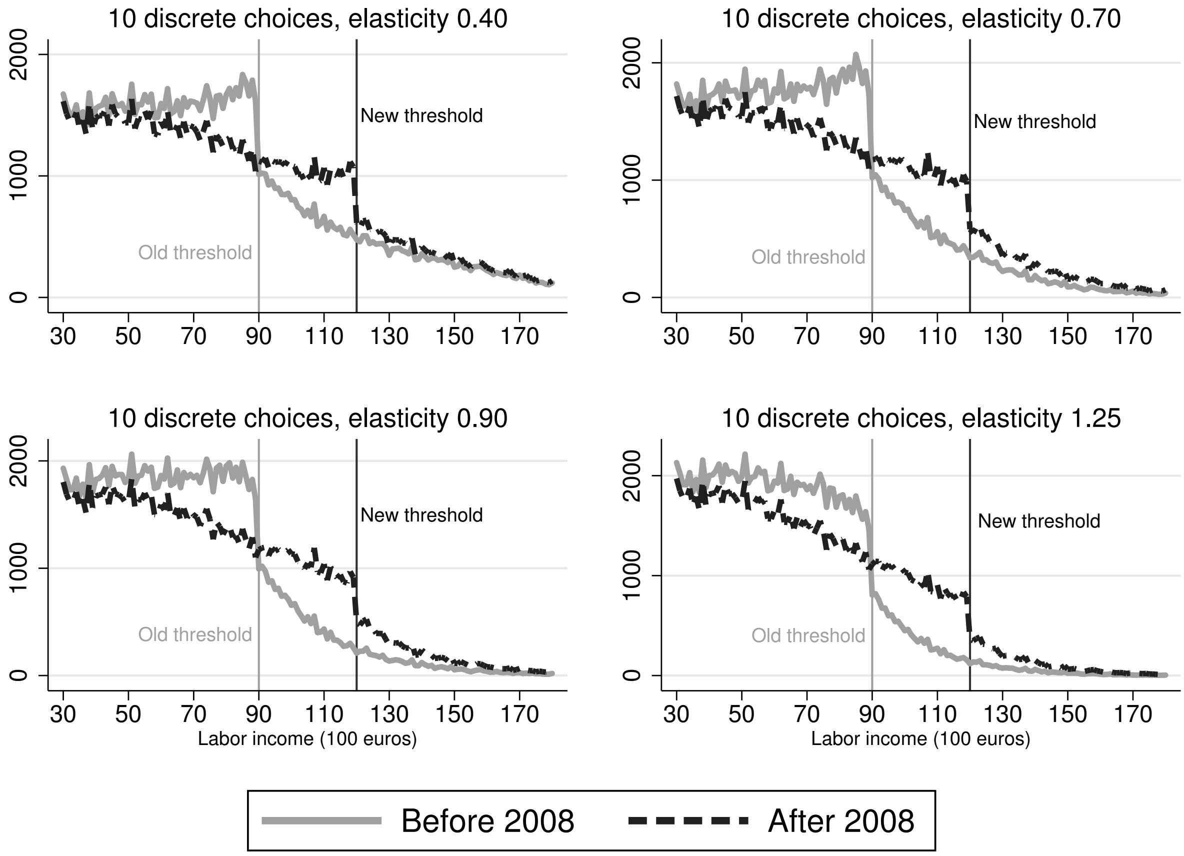

Notes: Figure presents simulated earnings distributions before (gray solid line) and after (black dashed line) the 2008 reform within an income range of 0–18,000 euros using different options for the elasticity parameter. The number of discrete earnings choices is set to 10. Figure shows that using different parameter values does not significantly affect the shape of the earnings distribution, but larger elasticity somewhat increases the broader changes in the distribution. Overall, Figures 8 and 9 illustrate that the number of discrete earnings choices appear to be more important in determining the shape and nature of responses to the relocation of the income threshold than the assumed underlying earnings elasticity parameter.

Figure 9: Simulated earnings distributions with different elasticity parameters

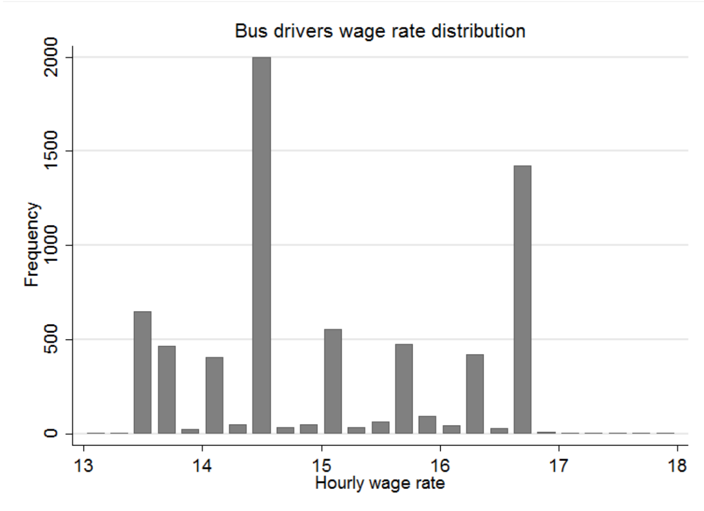

(a) Hourly wage rate distribution of bus drivers

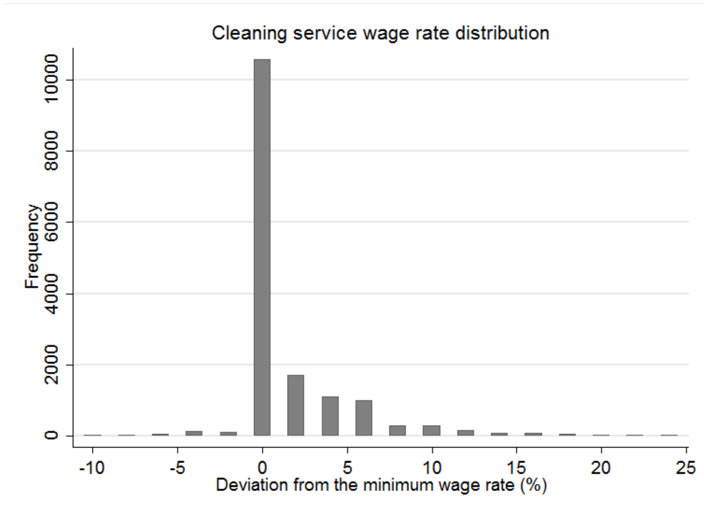

(b) Wage distribution relative to personal minimum wage, cleaning services

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 personal 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 10: Discrete wage rate opportunities for wage earners

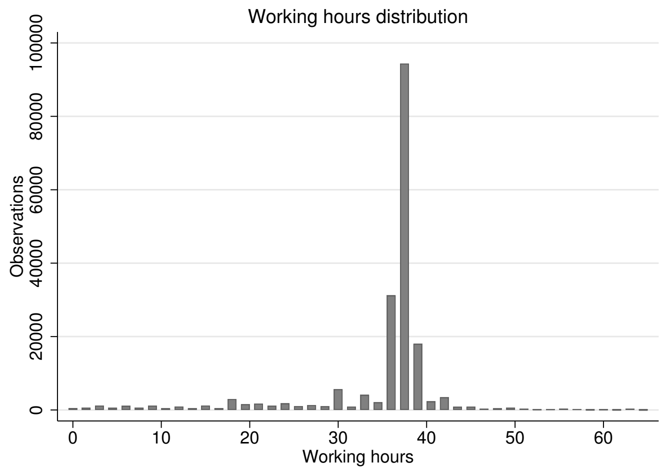

(a) Aggregate working hours distribution

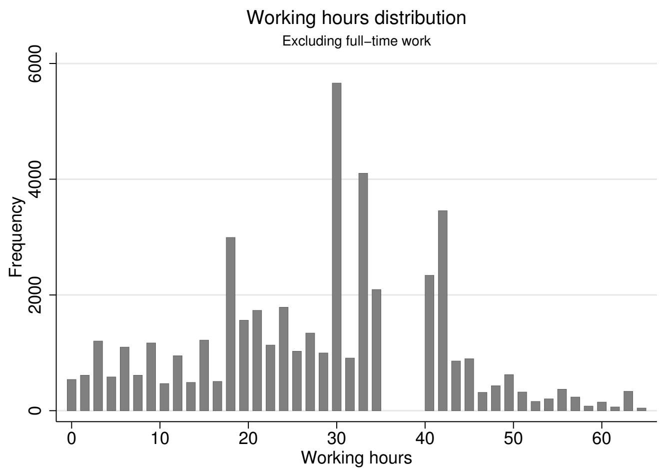

(b) Excluding typical full-time working hours

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 11: Aggregate working hours distributions for wage earners

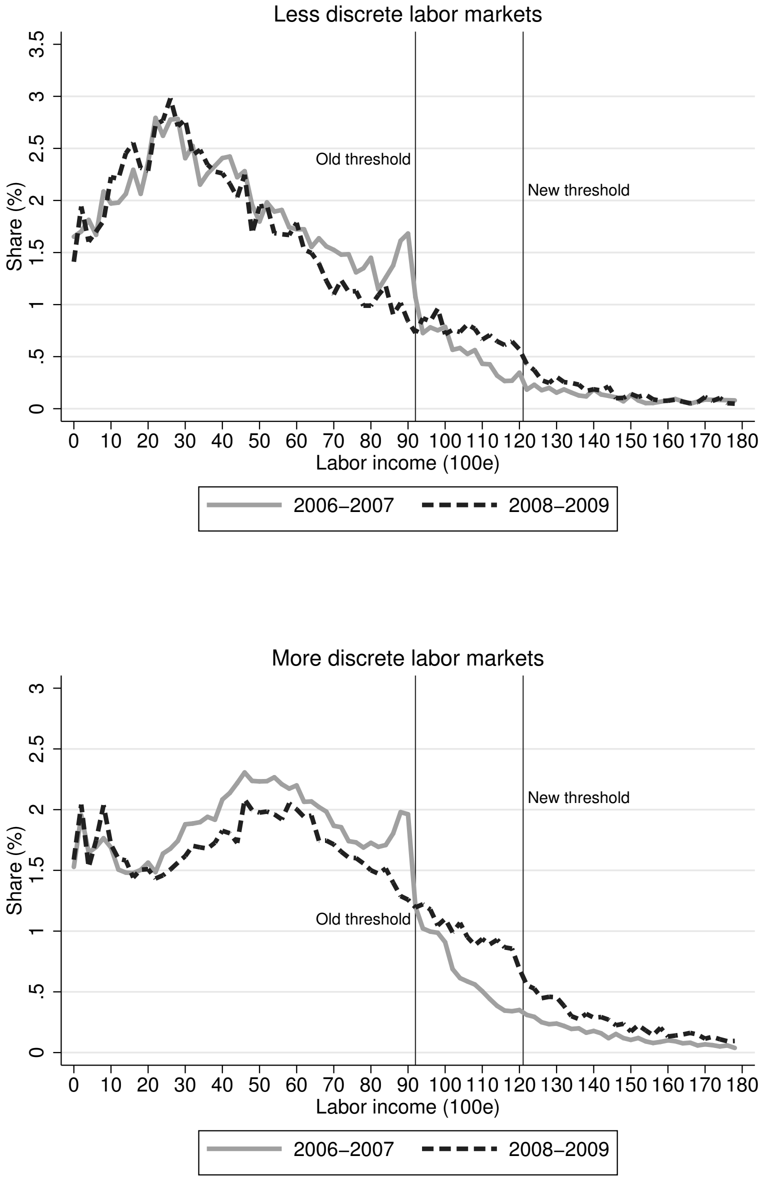

Notes: Figure presents the observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 200 euros for students with the default 9 subsidy months in each year working 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, and manufacturing and construction. 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 12: Labor income distributions before and after 2008 for students working in less discrete and more discrete labor markets

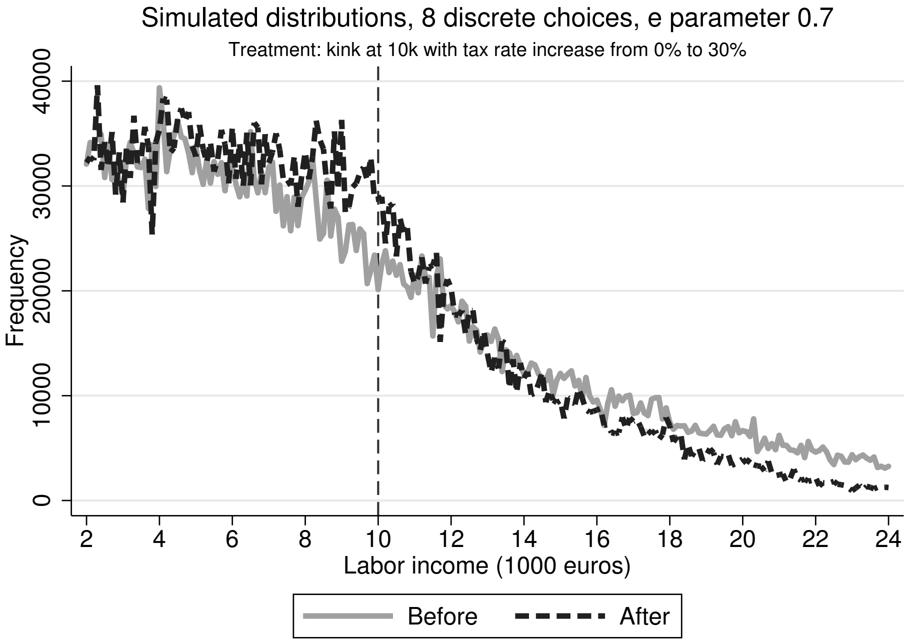

Notes: Figure presents the simulated earnings distributions before (gray solid line) and after (black dashed line) the introduction of a tax rate kink point at 10,000 euros. The marginal tax rate below the kink is 0% and 30% above it after the introduction of the kink. The underlying elasticity parameter of 0.7 and the assumption of 8 available earnings choices within 0–25,000 euros is used in the simulation in the figure. All individuals face an inflation of 2% in the after period.

Figure 13: Simulated income distributions with a 30 percentage-point tax increase at 10,000 euros in the after period

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 side 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	Administration & Publ. 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 side in the data.

Table 3: Parameter values in the simulation model

Parameter	Value
Marginal tax rate (τ)	
Below the notch	0.22
Above the notch	0.61
Size of the notch	500e
Virtual income (R)	
Before	4,100e
After	3,600e
Location of the notch (income threshold)	
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: Estimates for broader changes in the simulated earnings distributions using different model specifications

	Discrete choices				
	5	*10*	*15*	*20*	*30*
Elasticity: 0.4	2.53 (5.76)	4.51 (4.22)	5.44 (3.87)	5.77 (4.20)	5.41 (5.01)
Elasticity: 0.5	3.07 (5.72)	5.38 (4.06)	6.46 (3.80)	**6.80 (4.00)**	6.36 (4.76)
Elasticity: 0.6	3.56 (5.69)	6.13 (4.12)	**7.30 (4.00)**	**7.65 (4.07)**	**7.12 (4.42)**
Elasticity: 0.7	3.97 (5.76)	**6.74 (3.93)**	**7.99 (3.81)**	8.34 (4.11)	**7.72 (4.72)**
Elasticity: 0.8	4.32 (5.47)	**7.25 (4.08)**	8.55 (3.78)	8.89 (3.78)	**8.20 (4.54)**
Elasticity: 0.9	4.61 (5.23)	**7.64 (3.98)**	8.99 (3.78)	9.35 (4.06)	8.60 (4.32)
Elasticity: 1.10	4.99 (4.80)	8.10 (3.84)	9.51 (3.66)	9.91 (3.72)	9.13 (4.15)
Elasticity: 1.25	5.11 (5.05)	8.16 (3.93)	9.54 (3.91)	9.91 (3.71)	9.14 (3.88)

Notes: Table collects the simulated estimates and standard errors for the broader changes in the earnings distribution after the 2008 reform using different assumptions on the number of available discrete earnings choices and the elasticity parameter. An income interval of 0–7,700 euros is used in the estimation, capturing potential changes in the shape of the distribution below the local bunching region, i.e. 1,500 euros below the threshold. In our baseline empirical case in Figure 3, the estimate for broader changes within this income interval is 7.41(0.86). In the table, the highlighted estimates fall within a one standard deviation of the estimated broader income change in the observed distribution. The de-emphasized lines and columns represent the simulations in which the overall shape of the simulated earnings distributions do not qualitatively match the empirical distributions.

Table 5: Differences-in-differences estimates using a hypothetical tax reform and simulated data: Baseline results

	(1)	(2)	(3)	(4)
Elasticity parameter	*0.7*	*0.7*	*0.7*	*0.4*
No. of discrete choices	*10000*	*10*	*5*	*10*
D ln(1-tax rate)	**0.700***	**0.206***	**0.044***	**0.118***
Standard error	(1.05e-06)	(0.0004)	(0.0004)	(0.0003)
N	8,925,162	9,524,904	9,566,114	9,459,265
R^2	1.000	0.026	0.001	0.015

Notes: Table collects the simulated differences-in-differences estimates with different assumptions on the underlying elasticity parameter and the number of available discrete earnings choices. Using a large number of available choices (1,000) returns the same differences-in-differences estimate with the underlying disutility parameter in the model (column (1)). Limiting the number of available earnings choices significantly reduces the estimate compared to the this baseline case (columns (2)–(3)). In addition, limiting the number of available earnings choices reduces the R^2 statistic to similar levels observed in many empirical studies.

Table 6: Differences-in-differences estimates using a hypothetical tax reform and simulated data: Additional results

	(1)	(2)	(3)	(4)
Elasticity parameter	*0.7*	*0.7*	*0.7*	*0.7*
No. of discrete choices	*200*	*100*	*50*	*30*
D ln(1-tax rate)	**0.677*** **	**0.645*** **	**0.586*** **	**0.506*** **
Standard error	(8.70e-05)	(0.000147)	(0.000239)	(0.000330)
N	9,053,342	9,053,342	9,332,609	9,405,594
R^2	0.967	0.967	0.668	0.401

(a) Larger number of available discrete choices

	(1)	(2)	(3)
Elasticity parameter	*0.7*	*0.7*	*0.4*
No. of discrete choices	*10*	*5*	*10*
D ln(1-tax rate)	**0.183*** **	**0.0284*** **	**0.107*** **
Standard error	(0.000795)	(0.000730)	(0.000617)
N	9,551,174	9,575,844	9,475,453
R^2	0.006	0.000	0.004

(b) Smaller tax reform (from 0 to 10%)

Notes: Sub-table (a) collects the simulated elasticity estimates when using the underlying elasticity parameter of 0.7 and varying the available discrete earnings choices from 200 to 30. The results show that the observed elasticity estimates begin the significantly reduce only after including a sufficiently extensive discrete choice constraint. Sub-table (b) presents the baseline simulation estimates when using a smaller tax reform where the tax rate is increased from 0 to 10% (baseline: 30%) above the 10,000 euro kink point for the treatment group. Using a different-sized tax rate variation does not have an economically significant impact on the observed elasticity estimates.

Appendix A

Figures

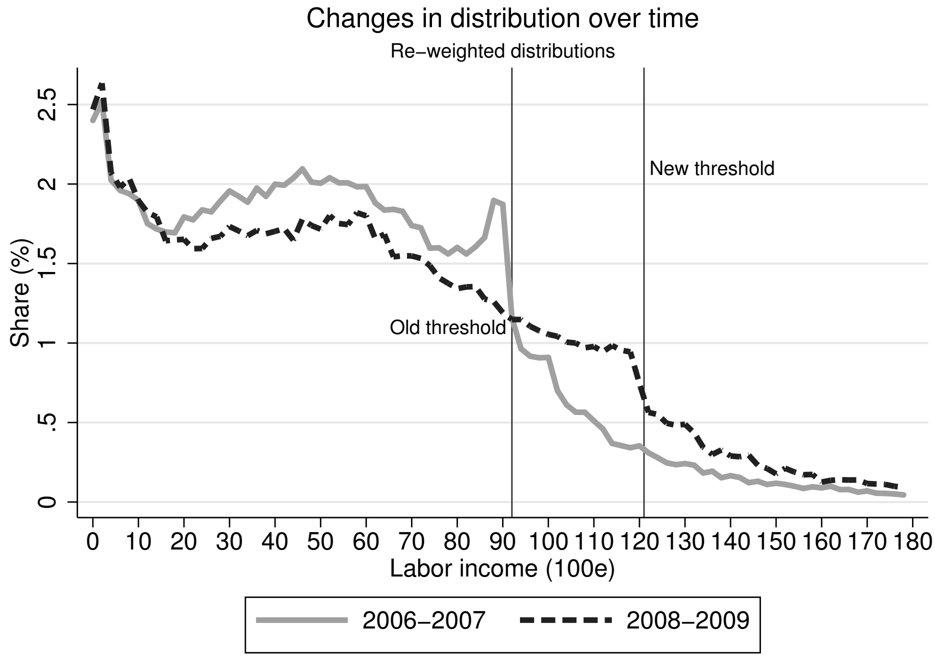

Notes: Figure presents the re-weighted observed relative earnings distributions before the reform in 2006–2007 (gray solid line) and after the reform in 2008–2009 (black dashed line) within an income range of 0–18,000 euros in bins of 200 euros for students with the default 9 subsidy months in each year. 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.

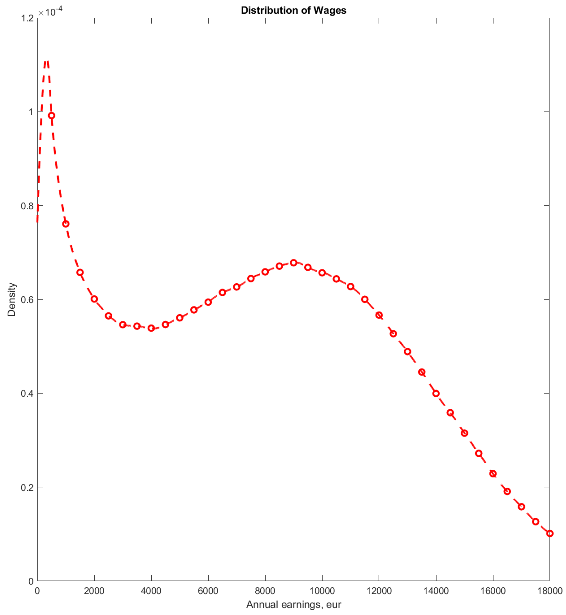

Notes: Figure presents the underlying earnings distribution used in the simulation model. The distribution is a combination of a power distribution and a normal distribution, which delivers an approximate match for the shape of the empirical earnings distribution of students in our empirical analysis. 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

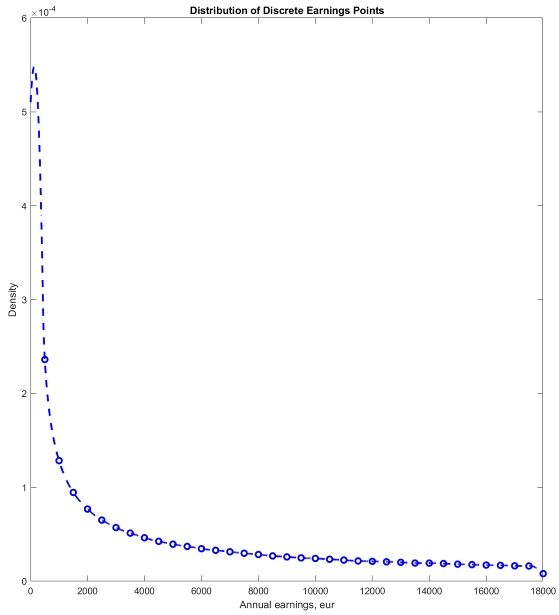

Notes: Figure presents the underlying probability distribution of discrete earnings choices utilized in the discrete choice model simulations. The large mass in the probability distribution at small earnings ensures that each individual has at least one available choice that produces positive utility 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

	Before 2008 (academic year 2006/2007)		**After 2008 (academic year 2008/2009)**	
Study subsidy months	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: Descriptive statistics, non-students aged 19–50, 1999–2013

	Individual characteristics				
	Age	Female	Labor income	Labor income > 500	Share with one employer
Mean	36.4	.48	25,912	.60	.81
Median	37	0	24,152	1	1
sd	8.914	.499	29,241	.49	.39
N	29,261,269	29,261,269	24,634,474	39,206,269	31,383,598
	Field of industry				
	Industry	Services	Admin. & Publ. Sector	Other/missing	
Mean	.27	.13	.42	.17	
sd	.44	.34	.49	.37	
N	39,206,521	39,206,521	39,206,521	39,206,521	

Notes: Table presents the descriptive statistics for all non-students in the data in 1999–2013. 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 side in the data.

Table A3: Income transition matrices, 2007–2008 and 2006–2007

2007–2008

		Year 2008					
		0–3,000	3,001–6,000	6,001–9,000	9,001–12,000	12,001–15,000	15,001–18,000
Year 2007	0–3,000	**0,38**	0,33	0,18	0,08	0,03	0,01
	3,001–6,000	0,11	**0,41**	0,34	***0,15***	0,04	0,01
	6,001–9,000	0,05	0,13	**0,36**	***0,33***	***0,10***	0,03
	9,001–12,000	0,03	0,06	0,19	**0,39**	***0,24***	0,09
	12,001–15,000	0,03	0,05	0,12	0,22	**0,32**	***0,26***
	15,001–18,000	0,05	0,07	0,13	0,20	0,23	**0,32**

2006–2007

		Year 2007					
		0–3,000	3,001–6,000	6,001–9,000	9,001–12,000	12,001–15,000	15,001–18,000
Year 2006	0–3,000	**0,37**	0,35	0,18	0,07	0,02	0,01
	3,001–6,000	0,11	**0,38**	0,35	0,12	0,03	0,01
	6,001–9,000	0,05	0,15	**0,45**	0,26	0,07	0,02
	9,001–12,000	0,04	0,07	0,26	**0,38**	0,18	0,07
	12,001–15,000	0,03	0,06	0,12	0,27	**0,35**	0,18
	15,001–18,000	0,07	0,06	0,10	0,16	0,27	**0,36**

Notes: Table presents the income transition matrices for 2007–2008 (upper table) and 2006–2007 (lower table). The matrices describe how students with 9 subsidy months changed their earnings within each 3,000 euro base-year earnings bin. For example, the first column of the first row shows the share of students earnings 0–3,000 euros in both of the years, and the next column shows the share of students who earned 0–3,000 in the first year and 3,001–6,000 in the next year etc. The highlighted numbers in the upper table denote the cells where the share of students has increased by at least 3 percentage points when comparing 2007–2008 to 2006–2007, showing that large earnings increases exceeding the next earnings bin cutoff were more common for income bins below the old income threshold (9,200 euros).

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:

$$c_j = \sum_{i=0}^{p} \beta_i (z_j)^i + \sum_{i=z_L}^{z_H} \eta_i \cdot \mathbf{1}(z_j = i) + \varepsilon_j \tag{8}$$

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:

$$\hat{b}(z^*) = \frac{\sum_{i=z_L}^{z^*} (c_j - \hat{c}_j)}{\sum_{i=z_L}^{z^*} \hat{c}_j / N_j} \tag{9}$$

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, $\hat{b}_M(z^*) = (\sum_{i=z>z^*}^{z_H} \hat{c}_j - c_j)$, 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 $z^* + \triangle z$, 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 (8) 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 +/- 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.

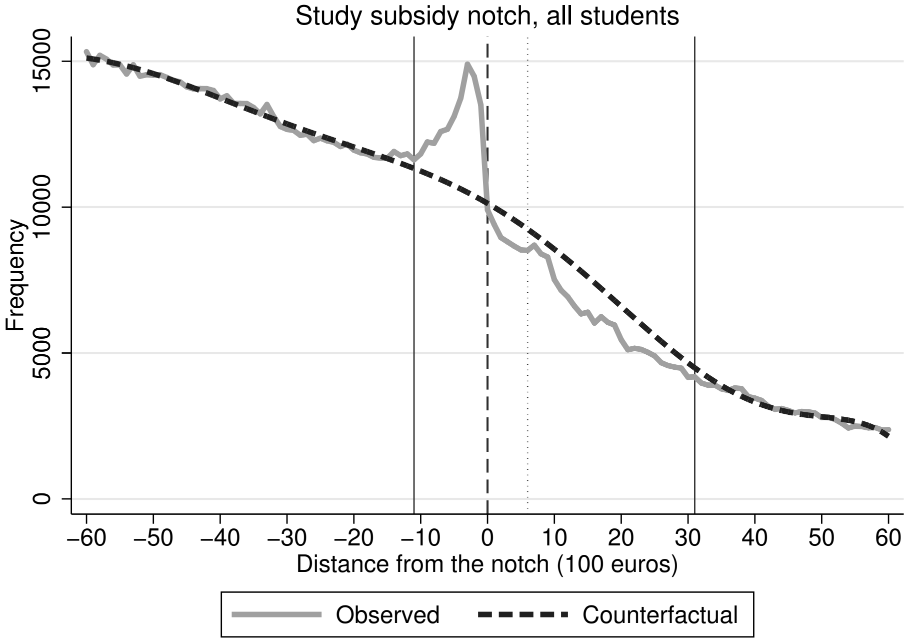

Notes: Figure presents the observed earnings distribution (gray solid line) and the estimated counterfactual distribution (black dashed line) around the income threshold (denoted by zero in the figure) in bins of 100 euros for all students using pooled data from 1999–2013. The first and second solid 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

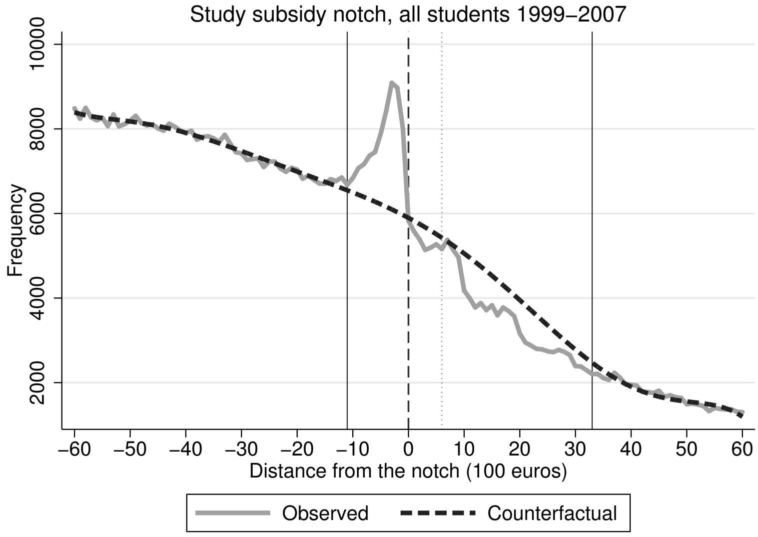

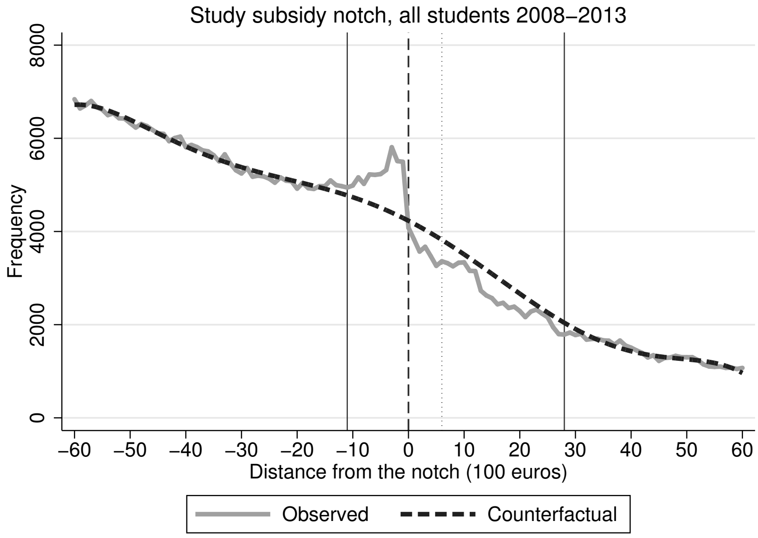

Notes: Figure presents the observed earnings distributions (gray solid line) and the estimated counterfactual distributions (black dashed line) around the income threshold (denoted by zero in the figure) in bins of 100 euros for all students before (1999–2007) and after (2008–2013) the 2008 threshold 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

$$e(z^*) \approx (z^H/z^*)^2/(\triangle t/(1-t)) \tag{10}$$

where $(1-t)$ is the net-of-tax rate at the notch, and $\triangle t$ 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 5.4 in the main text.[19]

[19]Furthermore, 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.