Essays on Tax Evasion and Enforcement and Intergenerational Wealth Mobility

A PhD dissertation by

Simon Halphen Boserup

presented at the

Department of Economics
Faculty of Social Sciences, University of Copenhagen

March 2013

Contents

Summary

This dissertation is comprised of three self-contained papers in two separate fields of economics. The first two chapters are closely related. Both are empirical applications of a structural model of tax evasion and enforcement. The model is applied to Danish, administrative tax return and enforcement data obtained from a large-scale experiment carried out by the Danish tax authorities, SKAT, in collaboration with researchers at University of Copenhagen and UC Berkeley. They investigate the redistributive implications of tax evasion and enforcement and the relative efficacy of different instruments available to the tax authorities in combating tax evasion. The last chapter contributes to the literature on intergenerational mobility by studying the correlation of wealth across generations in a Danish context using administrative wealth records for three generations of Danes in the years 1983-2011.

Chapter 1, *"Tax Evasion, Information Reporting, and the Regressive Bias Prediction"* (joint with Jori Veng Pinje), investigates a prediction from the tax evasion literature that optimal auditing induces a regressive bias in effective average tax rates compared to statutory rates, reducing the degree of redistribution in the tax system. We show that a model building on both rational tax evasion and the well-established fact that some taxpayers are inherently honest, and which takes into account that modern tax collection relies on information reporting from third parties in addition to traditional auditing, can convincingly replicate the moments and correlations of tax evasion and probabilities of audit. We find that both reduced-form evidence and simulations are in accordance with the prediction of regressive bias when conditioning on information reporting. However, the use of information reporting counteracts the regressive bias generated by optimal evasion and auditing behavior and, as a consequence, the bias vanishes when considering the degree of redistribution in the overall economy.

Chapter 2, *"A Structurally Estimated Model of Tax Evasion and Enforcement"* (joint with Jori Veng Pinje), estimates a structural a model of tax evasion and optimal auditing, inspired by the empirical congruity of the model developed in Chapter 1 to Danish data, and assesses the relative efficacy of instruments to deter tax evasion. We find that the policy instruments that work along the intensive margin of tax evasion (audits and penalty rates) are less effective in combating tax evasion than instruments working along the extensive margin of tax evasion (third-party information reporting and the share of honest taxpayers in the population).

Chapter 3, *"Intergenerational Wealth Mobility: Evidence from Danish Wealth Records of Three Generations"* (joint with Wojciech Kopczuk and Claus Thustrup Kreiner), provides empirical evidence on the intergenerational mobility of wealth using Danish, administrative wealth records, linking data for three generations. The preferred estimate of the intergenerational wealth elasticity (IWE) is 0.2 (and 0.27 when limiting attention to those with positive wealth). We construct a theoretical framework that allows for understanding the variability of the IWE across time, samples, and countries. It highlights that the IWE can be interpreted as the weighted average of elasticities corresponding to different sources of intergenerational correlation that may in principle vary in importance across different contexts. Surprisingly, the IWE is found empirically to be very stable when estimated for different age groups, when using parents-grandparents pairs instead of children and parents, when eliminating bequests, and when explicitly shutting down many of the potential channels behind intergenerational wealth mobility, including income and education. This suggests that parental wealth is a sufficient statistic for the channels that we control for and those that vary across different samples, that is, the effect of these parental characteristics on wealth of children can be summarized by their effect on wealth of parents. Finally, exploiting information for three generations, we find that the standard child-parents elasticity severely underestimates the long-term persistence in the formation of wealth across generations. In particular, we find that either the true elasticity is significantly underestimated or that grandparental characteristics matter beyond information incorporated in parental characteristics. We also find evidence supporting the presence of a persistent dynastic component, implying that different families will gravitate towards different wealth levels over generations thereby limiting the intergenerational wealth mobility.

Resumé (Summary in Danish)

Denne afhandling består af tre selvstændige artikler inden for to forskellige områder af den økonomiske litteratur. De første to kapitler er nært beslægtede. Begge er empiriske anvendelser af en strukturel model for skatteunddragelse af personlig indkomst og skattemyndighedernes forsøg på at stække denne. Modellen anvendes på danske, administrative data fra et storstilet eksperiment udført af SKAT i samarbejde med forskere ved Københavns Universitet og UC Berkeley. Kapitlerne studerer de omfordelingsmæssige konsekvenser af skatteunddragelse og skattemyndighedernes kontrolindsats, samt den relative effektivitet af forskellige instrumenter til rådighed for skattemyndighederne til bekæmpelse af skattesnyd. Det sidste kapitel bidrager til litteraturen om intergenerationel mobilitet ved at studere sammenhængen mellem formuer på tværs af generationer i en dansk sammenhæng. Dette studie baserer sig på registerdata om formuer mm. for tre generationer af danskere i årene 1983-2011.

Kapitel 1, *"Tax Evasion, Information Reporting, and the Regressive Bias Prediction"* (skrevet i samarbejde med Jori Veng Pinje), undersøger en forudsigelse fra skatteunddragelseslitteraturen, om at der vil forekomme regressiv bias i effektive gennemsnitlige skattesatser sammenlignet med lovbestemte skattesatser, når skattemyndighederne handler optimalt. Vi viser, at en model der bygger på både rationel skatteunddragelse og det veletablerede faktum, at nogle skatteydere i sagens natur er ærlige, og som tillige tager højde for, at moderne skatteopkrævning i høj grad bruger tredjepartsrapporterede indkomstoplysninger, overbevisende kan genskabe de primære momenter og korrelationer i data. Vi finder, at både reduceret-form resultater og simuleringer er i overensstemmelse med forudsigelsen om regressivt bias, betinget på graden af tredjepartsrapportet indkomst. Dog modvirker tredjepartsrapporteret indkomst det regressive bias, således at det forsvinder når vi betrager økonomien som helhed.

Kapitel 2, *"A Structurally Estimated Model of Tax Evasion and Enforcement"* (skrevet i samarbejde med Jori Veng Pinje), estimerer en strukturel model for skatteunddragelse og optimal kontrol inspireret af den udmærkede empiriske overenstemmelse mellem danske data og modellen, der blev udviklet i kapitel 1. Vi bruger modellen til at vurdere den relative effektivitet af skattemyndighedernes forskellige instrumenter til bekæmpelse af skatteunddragelse. Vi finder, at de instrumenter, der arbejder langs den intensive margin for skatteunddragelse (kontroller og bøder) er mindre effektive i bekæmpelsen af skatteunddragelse end instrumenter, der arbejder langs den ekstensive margin af skatteunddragelse (tredjepartsrapporteret indkomstsinformation og andelen af ærlige skatteydere i befolkningen).

Kapitel 3, *"Intergenerational Wealth Mobility: Evidence from Danish Wealth Records*

of Three Generations" (skrevet i samarbejde med Wojciech Kopczuk og Claus Thustrup Kreiner), studerer empirisk den intergenerationelle mobilitet målt ved formuer. Vi anvender danske registerdata for formuer for tre generationer af danskere. Det foretrukne estimat for den intergenerationelle formueelasticitet (IFE) er 0,2 (og 0,27 når vi begrænser os til kun at betragte personer med positiv formue). Kapitlet opstiller desuden en teoretisk ramme, der giver mulighed for at forstå variationen i IFE over tid, på tværs af forskellige stikprøver og på tværs af lande. Det fremhæves hvordan IFE kan fortolkes som et vægtet gennemsnit af elasticiteter svarende til de forskellige kilder til intergenerationel korrelation, som i princippet kan variere fra kontekst til kontekst. Overraskende viser det sig, at IFE er meget stabil på tværs af forskellige aldersgrupper, når man ser på forældre-bedsteforældre i stedet for børn og forældre, når vi udelukker muligheden for arv til overførsel af formuer og når vi eksplicit kontrollerer for en række potentielle kilder til intergenerationel formuemobilitet, inklusive indkomst og uddannelsesniveau. Det tyder altså på, at forældrenes formue kan anses som en såkaldt "sufficient statistic" for de kilder til intergenerationel mobilitet, som vi kontrollerer for. Dvs. at effekten af forældrenes forskellige karakteristika på deres børns formue kan opsummeres ved størrelsen af forældrenes formue. Endelig finder vi ved at anvende data for alle tre generationer, at den gængse IFE målt mellem børn og forældre undervurderer betydeligt den langsigtede persistens i dannelsen af formue på tværs af generationer. Vi finder således at den sande IFE enten er groft undervurderet, eller at bedsteforældrenes karakteristika har en selvstændig betydning ud over forældrenes indflydelse på børnene. Vi finder ligeledes belæg for tilstedeværelsen af en såkaldt dynastisk "fixed effect", således at forskellige familier graviterer mod forskellige formueniveauer gennem generationer, hvilket modarbejder intergenerationel formuemobilitet.

Acknowledgements

This dissertation was written during my enrollment at the PhD program in Economics at the University of Copenhagen from 2008 through 2013.

During this process, a number of people deserve my thanks and grattitude. First of all, my advisor, professor Claus Thustrup Kreiner, for guideance, helpful discussions, critique, and inspirational coauthorship, and my former advisor, professor Peter Birch Sørensen, among other things for encouraging me to enroll in the PhD program. I would also like to thank my coauthor and fellow PhD student, Jori V. Pinje, for his companionship through thick and thin over the years, and coauthor Wojciech Kopczuk for great collaboration.

I also want to thank fellow PhD students (former, present, and visiting), faculty members, and others at the Department of Economics, who have made the department such an inspiring, motivating, and pleasant place to be.

In 2011 I had the wonderful opportunity to visit the Burch Center of Tax Policy and Public Finance at UC Berkeley during the spring semester. It was a tremendously rewarding stay for me, both academically and personally. I am especially grateful to professor Alan Auerbach for inviting me to visit the Burch Center, and to both him and professor Emmanuel Saez for taking the time to discuss my research with me and for giving me helpful critique and comments along the way. I gratefully acknowledge the financial support from the Euroclear stipend, Augustinus Fonden, and Oticon Fonden that made my stay at UC Berkeley possible.

Last, but not least, I am grateful to my wife, Lise, and my two sons for enduring the ups and downs, especially in the finishing stages of the writing process, and for following me abroad and participating in that adventure when that was where my work took me.

Simon Halphen Boserup
Copenhagen, March 2013

Chapter 1

Tax Evasion, Information Reporting, and the Regressive Bias Prediction*

Simon Halphen Boserup† and Jori Veng Pinje

Department of Economics, University of Copenhagen

November 2012

Abstract

A robust, but untested, prediction from the tax evasion literature is that optimal auditing induces a regressive bias in effective average tax rates compared to statutory rates, reducing the degree of redistribution in the tax system. Using Danish administrative data, we show that a calibrated structural model of rational tax evasion and tax enforcement can convincingly replicate the moments and correlations of tax evasion and probabilities of audit once we account for the presence of information reporting in the tax compliance game. We find that both reduced-form evidence and simulations are in accordance with the prediction of regressive bias when conditioning on information reporting. However, information reporting counteracts the regressive bias generated by optimal evasion and auditing behavior and, as a consequence, the bias vanishes when considering the degree of redistribution in the overall economy.

*We are grateful to the Danish tax administration, SKAT, for providing data access and to Peer Skov and the Rockwool Foundation Research Unit for providing data and assistance on black market activities. We thank Alan Auerbach, Jonathan Feinstein, Henrik Kleven, Claus Thustrup Kreiner, David Dreyer Lassen, Emmanuel Saez, Suzanne Scotchmer, and Joel Sobel as well as seminar participants at Nordic Workshop on Tax Policy and Public Economics, UC Berkeley, the 6th Royal Economic Society PhD Meeting at City University London and the EDGE Meeting at University College Dublin for helpful comments and suggestions.

†Corresponding author. Tel. +45 2649 2696. *E-mail address:* simon.h.boserup@econ.ku.dk. Department of Economics, University of Copenhagen, Øster Farimagsgade 5, building 26, DK-1353 København K, Denmark.

1 Introduction

In this article, we develop a structural model of tax evasion and enforcement in a population of taxpayers. Highly detailed Danish administrative data allows us to perform a meaningful calibration exercise to investigate the model's ability to explain tax evasion and the tax agency's enforcement strategy. We show that the model's predictions closely match key empirical relationships in the data and, in particular, we provide the first empirical evidence of the regressive bias prediction established in the theoretical literature on tax evasion and optimal enforcement (see for example Reinganum and Wilde, 1986; Cremer, Marchand, and Pestieau, 1990; Sanchez and Sobel, 1993; Erard and Feinstein, 1994).

The potential for tax evasion requires a distinction between the statutory tax system and the *effective* tax system. Tax evaders pay less taxes than they should and this implies a wedge between statutory and effective average tax rates. The regressive bias prediction states that this wedge is larger for high-income taxpayers than for low-income taxpayers – even when the enforcement regime is revenue maximizing. Thus, the tax system may be substantially less redistributive than intended by the tax code. As shown by Scotchmer (1992), the prediction of regressive bias is theoretically robust. Model variations in the literature consistently arrive at regressively biased effective average tax rates.

The intuition behind this prediction is the following: The tax compliance game played by the tax agency and taxpayers is a screening problem in which high-income taxpayers can increase their expected payoff by imitating low-income taxpayers. If not all taxpayers can be audited, the tax agency should optimally prioritize tax returns reporting low income. Rather than eliminating tax evasion altogether, budget-constrained optimal enforcement primarily discourages very low reports by high-income individuals. Due to the optimal regressivity in tax enforcement, evading taxes on the margin subjects a low-income taxpayer to a greater risk of getting caught than a high-income taxpayer, which tends to make high-income taxpayers evade more. In equilibrium, the decreasing relationship between the probability of audit and reported income and the increasing relationship between evaded taxes and true income lead to an increasing wedge between the statutory average tax rate and the effective average tax rate as a function of true income, i.e., a regressive bias. Figure 1(a) illustrates how the wedge between the effective average tax rate, τ^{eff}, and the average tax rate as implied by the statutory tax system, τ, is increasing in true income.

There is one important exception to the regressive bias result: when the tax agency uses *ex ante* observable population variables, such as gender, age, occupation, or employer-reported salaries, to predict true incomes, there may be no bias or even progressive bias in the population as a whole. How this plays out in particular economies will determine the appropriate way to account for the redistributional aspects of tax evasion and enforcement through economic policy. Scotchmer (1987) shows that when tax agencies facilitate prediction of taxpayers' true income by dividing taxpayers into *audit groups*, upon which the agency conditions its enforcement strategy, effective average tax rates remain regressively biased within audit groups but the direction of the bias between groups is ambiguous. The aggregate bias depends on the predictive power of the signals (i.e., the *ex ante* known population variables) and the allocation of audit resources across audit groups. Consequently, the regressive bias prediction should be interpreted as a *within-audit-group* phenomenon. Figure 1(b) illustrates the aggregate relationship between effective average tax rates, τ^{eff}, and true income, which is a composite of relationships within multiple audit groups, τ_i^{eff}. Whereas the regressive bias prediction remains valid within audit groups, effective tax rates may be *progressively biased* across audit groups.

The mechanism driving the result is that some low-income taxpayers benefit from being high-income individuals within their audit group while some high-income taxpayers instead are low-income taxpayers within their audit group. This reclassification changes the risk of being audited and, hence, the *ex ante* effective tax rate. In addition, the tax agency can more efficiently target high-income individuals by modifying the distribution of audit resources between audit groups. If the observable signal of true income is stronger or audits are more abundant among high-income taxpayers, progressive bias between groups may dominate in the aggregate.

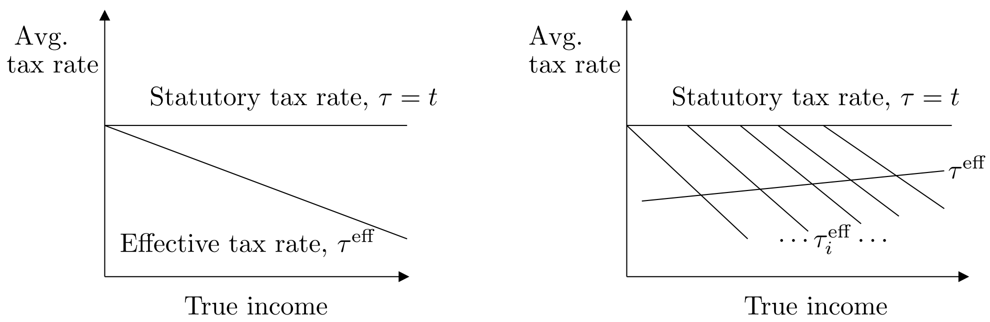

(a) The Regressive Bias Result. (b) Aggregation Across Audit Groups.

Figure 1. Correlation Structure of Effective Average Tax Rates.
Notes: τ is the statutory average tax rate (here, constant at $\tau = t$), τ_i^{eff} is the effective average tax rate within audit group i, and τ^{eff} is the aggregate effective average tax rate.

We apply a specific theoretical structure to Danish administrative data on tax compliance/evasion and show that the empirical properties of tax evasion, tax enforcement, and effective tax rates are convincingly replicated by a screening game between a tax agency and taxpayers. To this end, we combine insights from two main sources, Kleven, Knudsen, Kreiner, Pedersen, and Saez (2011) and Erard and Feinstein (1994). In the former, the authors collect a uniquely detailed micro-data set based on a random sample of Danish taxpayers containing pre- and post-audit incomes and taxes, as well as reports on income, proxies for audit probabilities, etc. They show that third-party reported income is by far the best predictor of true income compared to other population variables. Since the Danish tax agency, SKAT, does in fact use these information reports extensively in its enforcement efforts, they are ideal for constructing audit groups.[1] Based on this insight, we generalize Erard and Feinstein's within-audit-group model to describe tax evasion and optimal enforcement both within and between audit groups. We calculate an internally consistent set of model parameters directly from data and calibrate the tax agency's budget to match the simulated level of tax evasion to data. We evaluate the model numerically and find that applying structure to the data yields results in close correspondence with a minimal-assumptions reduced-form approach. This model convincingly replicates tax evasion behavior for both wage earners and the self-employed although these two groups

[1] Other recent papers demonstrate the importance of explicitly considering information reporting. Phillips (2010) demonstrates the predictive power of an indirect measure of third-party reported information in US data and Pomeranz (2010) demonstrates the general importance of information as a deterrent of VAT evasion in a sample of Chilean firms.

differ markedly in terms of the propensity to evade taxes and the extent and distribution of third-party reported income. We conclude that (statically optimized) tax evasion and tax enforcement is sufficient to generate the observed structure of effective average tax rates.

Overall, our micro-data on Danish taxpayers suggests that there is a regressive bias within audit groups. Between audit groups, tax rates are progressively biased to such an extent that tax rates are actually progressively biased in total income. Thus, our findings support the regressive bias prediction at the theoretical level but not as an aggregate empirical outcome in Danish data; specifically, our results correspond closely to the structure of effective tax rates conjectured in Scotchmer (1987). Moreover, using information about the enforcement regime, we find evidence suggesting that the actual audit regime exhibits the key qualitative features of an optimal audit regime and that the correlation structure of effective average tax rates is, indeed, caused by the theorized combination of optimal enforcement and tax evasion.

In model simulations, the covariance structure of effective average tax rates is robust to parameter variations. In view of this, we predict that similar empirical relationships would be found in data from any tax agency that employs, as does the Danish tax agency, a strong signal in predicting true incomes. The model also suggests that enforcement regimes employing information reporting to a lesser extent may be substantially more regressive.

Our results have important implications for policy. Due to the theoretical robustness of the regressive bias prediction, it has been argued (e.g., in Scotchmer, 1992) that governments could increase the progressivity of the income tax schedule to counter regressive bias inherent in optimal tax enforcement. However, our results imply that such a policy adjustment is undesirable. In the first place, adjusting tax rates cannot eliminate the inequity between taxpayers that evade taxes and taxpayers that do not. Secondly, there may be no regressive bias to correct in the aggregate due to tax agencies' use of third-party reported information in tax enforcement. If such is the case, the policy priority is correcting the horizontal inequity between evaders and non-evaders rather than the distortion of redistribution between high- and low-income taxpayers – for this purpose allocating more resources to the tax agency or collecting more information *ex ante* are superior approaches.

Our results illustrate the importance of including information reports in empirical analyses of tax evasion and enforcement. Neglecting to account for information reports may lead to counterintuitive comparative statics estimates such as for tax evasion with respect to total income or marginal taxes. This may partly explain the empirical literature's lack of consensus with respect to basic correlations between measures of tax evasion, tax rates, and income.[2]

We now proceed to the main body of the paper. Section 2 develops our model of the tax compliance/evasion game. Section 3 outlines the Danish tax system and describes the main features of the data. Section 4 describes the calibration of parameters, outlines the numerical strategy and establishes the correspondence of data and model-generated output. Section 5 concludes. The Appendix provides details of the numerical implementation and a description of black market activity in Denmark.

2 Theory: A Model of Income Tax Auditing Subject to Information Reporting

Several current theories are capable of analyzing behavior *within* audit groups, i.e., conditional on pre-defined groups based on *ex ante* observable information. However, as we wish to analyze *aggregate* reporting behavior as well as the tax agency's overall response, we need a model that can encompass a population of taxpayers, i.e., several audit groups. To this end we generalize the model in Erard and Feinstein (1994) to incorporate a population that is heterogeneous in third-party income reports.[3]

[2]For example, Feinstein (1991) finds a negative effect of marginal tax rates on underreporting, whereas Clotfelter (1983) finds a positive effect. With respect to the effect of income on underreporting, Feinstein (1991) finds no effect, whereas Clotfelter (1983) finds a positive effect.

[3]We use a different specification for penalties in case of detected evasion compared to Erard and Feinstein (1994). We model penalties as proportional to evaded taxes rather than evaded income as this is also the structure of the actual Danish penalty system.

Erard and Feinstein (1994) introduce noise in taxpayer reports by incorporating the stylized fact that some taxpayers report their incomes honestly, even when they have ample opportunity to evade taxes. This is also the case in our data as we demonstrate in Section 3. As argued in Erard and Feinstein (1994), including inherently honest taxpayers increases the realism and usefulness of the model: it eliminates several potential equilibria and leaves them with a unique revenue maximizing equilibrium prediction. Further, it eliminates the unrealistic feature of earlier models that the tax agency in equilibrium would know the true incomes of all taxpayers *before* the actual audit.[4] Thus, for each tax return filed by a particular taxpayer, the tax agency decides whether or not to audit based on the expected reports of dishonest and honest taxpayers and the likelihood that any particular tax return is fraudulent.

To develop a model that we can apply to data, we extend the model in Erard and Feinstein (1994) to account for the tax agency's use of information reports. As shown by Kleven et al. (2011), in the Danish context, third-party reported income is by far the most powerful predictor available, making it an ideal candidate for defining audit groups. However, as this variable, like true income, is intuitively best understood as a continuous variable, we allow the tax agency to choose audit functions contingent on the third-party information of a particular taxpayer and interpret each *level* of third-party reported income as an audit group. Reflecting the very low evasion rates on third-party reported income in our data, we use the simplifying assumptions that these reports are always correct and are common knowledge to both taxpayer and tax agency. Overall, the probability that a particular taxpayer is audited depends both on the exogenous signal, i.e., third-party reported income, and the endogenously determined reported income.

The structure of the model is illustrated in Figure 2. The tax agency selects the audit regime subject to a budget constraint without being able to commit to an audit strategy. The audit schedule for a particular audit group (i.e., conditional on a particular third-

[4]A limitation of the modeling framework is that it does not explain why some taxpayers choose to report honestly. However, the model is well-suited for analyzing the behavior of rational tax evaders *given* that some taxpayers are, in fact, honest. Moreover, it provides a relatively simple framework for analyzing optimal enforcement in the face of this behavior and subject to the informational asymmetries inherent in the tax enforcement/compliance game.

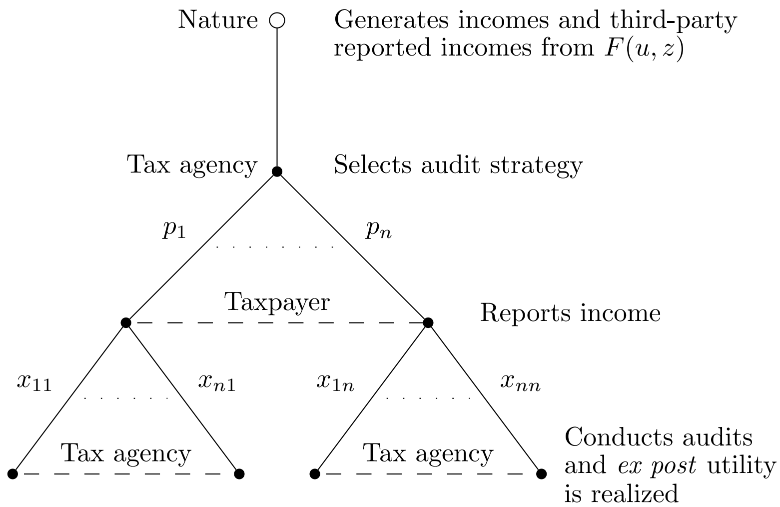

Figure 2. Game Tree.

party reported income level) is a function of taxpayers' reported *residual* incomes, i.e., income in excess of third-party reported income, reflecting our assumption that third-party reported income is common knowledge. The tax agency allocates its resources across different strata of the population so as to equalize the shadow values of extending resources to auditing taxpayers with different amounts of third-party reported income. Whereas the distribution of true incomes, conditional on information reports, is known, actual true incomes of individual taxpayers are private information. Taxpayers choose income reports subject to their expectations about the audit regime. Finally, the actual returns and the audit schedule are realized, audits are conducted, and tax revenue and *ex post* utilities, as measured by income net of taxes and any penalty payments, are realized.

2.1 Individual Reporting Behavior

Individual taxpayers have true taxable incomes y and report taxable incomes, $\tilde{y}$. Part of true income, z, is reported by third parties and is known to all parties. Therefore, $y = z + u$, where u is residual income, which can be positive or negative as it includes both, e.g., wages and deductions not reported by third parties. u is *ex ante* unknown and can only be ascertained by the tax agency by conducting a costly audit, which we assume reveals all of "true" residual income.[5] We denote the reported residual x, such

[5]We follow Erard and Feinstein (1994) in assuming that taxpayers do not incur a cost from filing taxes (time costs, hiring of a tax accountant, concealment costs etc.). Such costs have welfare consequences in the form of deadweight losses. Cremer and Gahvari (1994) show that a concealment technology that

that $x = \tilde{y} - z$.

Erard and Feinstein (1994) split taxpayers into two broad groups, honest and dishonest taxpayers, and assume that these two types differ only in reporting behavior, and that honesty is uncorrelated with true income. However, empirically the ratio of compliant to noncompliant taxpayers is not constant on the domain of u due to a large mass of correct reports around $u = 0$. The reason for this is that third-party reported income is such a strong signal of true income that, for many taxpayers, it is, indeed, virtually a perfect signal.[6] However, this can be remedied by a minimal departure from the assumptions of Erard and Feinstein (1994) by letting the ratio of honest to dishonest taxpayers differ on the domain of u. We define the densities of true income conditional on third-party reports $f^h_{u|z}$ and $f^d_{u|z}$ for honest and dishonest taxpayers, respectively. In addition, we define the total density function as $f_{u|z} = f^h_{u|z} + f^d_{u|z}$ and $F_{u|z}$ the conditional distribution function associated with $f_{u|z}$.

We follow Erard and Feinstein (1994) in assuming that taxes are linear in income.[7] Whereas honest taxpayers always report $x = u$, we assume that dishonest taxpayers are risk neutral and maximize expected utility given by expected income net of taxes and penalties

$$(1-t)\,z + p\,(x|z)\,[(1-t)\,u - \theta t\,(u-x)] + (1 - p\,(x|z))\,[u - tx]\,,$$

where t is the tax rate, θ is the penalty rate on tax evasion, and $p(x|z)$ is the audit

allows taxpayers to lower the probability of detection at a cost can affect the effective progressivity of the tax system. This may result in more or less progressivity depending on the exact specification of the concealment technology. However, their model assumes a constant audit probability, whereas our model implies a non-increasing audit probability on the domain of reports of dishonest taxpayers. In any case, whether or not such costs are important, our results in Section 4 indicate that they are not necessary to explain the correlation structure of effective average tax rates.

[6] In principle, such taxpayers could still evade taxes by claiming unwarranted deductions. This type of reporting behavior is virtually non-existent in our data. A possible explanation is that the burden of proof is on the taxpayer in such cases. On the other hand, having negative residual income (i.e., some deductions not subject to third-party reporting) allows for tax evasion by overstating the value of otherwise legal deductions.

[7] Clearly, this an abstraction but not an extreme one. Although the income tax schedule has three brackets, the average tax rates are much smoother. It would also be possible to perform the analyses using a full, nonlinear specification of taxes. We do not expect that the conclusions of this paper would be substantially affected by this change. Moreover, to accomodate the progressiveness of marginal income taxes as much as possible, in the empirial application of the model we allow the model's constant marginal tax rate to vary in z.

probability for report x given the level of third-party reporting z. The correct amount of taxes are paid with certainty on income reported by third parties, whereas taxes (and penalties) paid on residual income depends on both a taxpayer's evasion behavior and whether or not the taxpayer is audited.

In optimum, the taxpayer's choice must satisfy the first order condition

$$u = x + \frac{p(x|z) - \frac{1}{1+\theta}}{p'(x|z)}. \tag{1}$$

It is clear from Equation (1) that for $p(\cdot) = \frac{1}{1+\theta}$, $x = u$ and evasion is discouraged completely. However, $p \geq \frac{1}{1+\theta}$ is not compatible with equilibrium when the tax agency cannot commit to the audit regime: if evasion were completely discouraged, the tax agency would lower p for some x as a cost saving measure. Thus, in equilibrium $p(\cdot) \in \left[0, \frac{1}{1+\theta}\right)$. Furthermore, the incentive compatibility constraints on the tax agency's optimization problem implies that audit functions are decreasing on the domain of income reports (see Erard and Feinstein (1994) for a detailed demonstration of this point).

Given that $p'(x|z)$ is negative and $p(x|z) < \frac{1}{1+\theta}$, increasing the audit probability will, *ceteris paribus*, lower tax evasion as the risk of getting caught is higher. Lowering $p'(x|z)$ (increasing its absolute value) also reduces tax evasion by increasing the risk of audit from taxes evaded on the margin.[8]

2.2 Optimal Audit Response

The tax agency chooses a continuum of audit schedules $p(x|z)$ and a budget allocation $B(z)$ for all z. In this way, the informational aspect of using third-party reported incomes to predict true income is incorporated into the population-wide equilibrium.[9] The audit

[8] Taxpayers' income returns must also satisfy the second order condition, $p''(x|z)(x-u) + 2p'(x|z) \leq 0$.

[9] In principle, the tax agency could also condition audit schedules on other population variables such as gender, age, occupation, etc. However, as Kleven et al. (2011) show, these variables are less powerful as predictors. Conditioning on whether the taxpayer was audited in previous years would complicate matters as it would introduce a dynamic aspect to reporting decisions. However, as observations on past audits are not employed in SKAT's actual audit scheme, this limitation is unlikely to affect the fit of our model. In addition, the statute of limitations for retrospective audits is limited to 14 months.

schedule is chosen to maximize expected revenue (taxes plus fines)[10]

$$\int \left(\int_{\underline{x}}^{\bar{u}} \left[p\left(x|z\right) \left(t\mathrm{E}\left(y|x,z\right) + \theta t \left(\mathrm{E}\left(y|x,z\right) - \tilde{y} \right) \right) + \left(1 - p\left(x|z\right)\right) t\tilde{y} \right] \mathrm{d}F_{x|z} \right) \mathrm{d}F_z$$

subject to the budget constraint

$$c \int \left(\int_{\underline{x}}^{\bar{u}} p\left(x|z\right) \mathrm{d}F_{x|z} \right) \mathrm{d}F_z \leq \int B\left(z\right) \mathrm{d}F_z \equiv \mathbf{B}, \tag{2}$$

where $F_{x|z}$ is the induced conditional distribution function for reported residual income, x, given third-party reported income, z; F_z is the marginal distribution function for z; and $B\left(z\right)$ is the proportion or density of the overall audit budget, $\mathbf{B}$, allocated to income reports with third-party reported income, z. For each (x, z), the tax agency must choose p to solve

$$\begin{aligned} \max_{p} \left\{ p \left[t\mathrm{E}\left(y|x,z\right) + \theta t \left(\mathrm{E}\left(y|x,z\right) - \tilde{y} \right) \right] + (1-p) t\tilde{y} \right\} \mathrm{d}F_{x|z} \mathrm{d}F_z \\ -\lambda\left(z\right) c \left[p \ \mathrm{d}F_{x|z} - B\left(z\right) \right] \mathrm{d}F_z, \end{aligned}$$

where $\lambda\left(z\right)$ is the Langrangian multiplier on the budget constraint. This implies a point-wise first order condition

$$tE(y|x,z) + \theta t E(y|x,z) - \theta t\tilde{y} - t\tilde{y} - \lambda\left(z\right) c \gtreqless 0, \tag{3}$$

which is greater than, equal to, or less than zero as $p = \frac{1}{1+\theta}$, $p \in \left(0, \frac{1}{1+\theta}\right)$, or $p = 0$. We look for equilibria in which the tax agency chooses a mixed strategy such that (3) holds with equality.[11]

As mentioned, our model is a generalization of the model in Erard and Feinstein (1994). Specifically, our model simplifies to theirs if $i)$ z is zero for all individuals, such

[10] Scotchmer (1992) shows that maximizing some measure of social welfare instead of expected revenue does not change the qualitative prediction that (within an audit group) there will be regressive bias, although it may change the distribution of resources across audit groups. The similarity of the observed and simulated distribution of resources, cf. Section 4, suggests that revenue maximation is not an inappropriate simplification in this context.

[11] The second order condition is $\frac{\partial \mathrm{E}(y|x,z)}{\partial p(x|z)} \geq 0$. In our simulations the solutions always satisfy this criterion.

that $F_{u|z} = F_u = F_y$, and ii) the ratio of honest to dishonest taxpayers, $\frac{f_u^h(u)}{f_u^d(u)}$, is constant on $[\underline{u}, \overline{u}]$. In this case, the problem becomes that of a partial optimization for a fixed $B(z)$ within an audit group. In this simpler version of the model, Erard and Feinstein (1994) show that the equilibrium audit and evasion functions have a number of useful properties. Due to the incentive constraints on reporting for high-income taxpayers, the audit function $p(x|z)$ is decreasing and continuous in reported income. The reporting function, $x(u|z)$ is strictly increasing in an upper region of the income domain and constant in a lower region as some taxpayers pool at the lowest possible report. As the audit and reporting functions are continuous and differentiable on the interior of the reporting domain, it is possible to solve for the equilibrium using methods of differential equations. In addition, as pooling occurs only at the lowest report, where the differential equation is undefined, sufficient conditions for equilibrium can be obtained by checking that the solution to the differential equation also satisfies the tax agency's first order condition for the lowest report, equivalent to (5) below. In the same way, we can leverage these properties to solve for the population-wide equilibrium as a range of within-audit-group equilibria coupled with the optimal budget distribution, $B(z)$.

The unique revenue maximizing equilibrium of the model is described by the collection of functions, $u(x|z)$ and $p(x|z)$, and the budget distribution, $B(z)$. Once $p(x|z)$ is determined, $u(x|z)$ is implicitly defined as the solution to the taxpayers' first order condition, and the tax agency chooses $p(x|z)$ such that (3) holds with equality. The two equations are connected by the tax agency's conditional expectation of taxpayers' true income given the reported income and third-party reports, $\mathrm{E}(y|x,z)$, which is

$$\mathrm{E}(y|x,z) = z + \frac{f_{u|z}^h(x)\,x + f_{u|z}^d(u(x|z))\,\frac{\partial u(x|z)}{\partial x}u(x)}{f_{u|z}^h(x) + f_{u|z}^d(u(x|z))\,\frac{\partial u(x|z)}{\partial x}}, \tag{4}$$

where the derivative $\frac{\partial u(x|z)}{\partial x}$ is derived from (1) by differentiating implicitly to get $\frac{\partial u}{\partial x} = 2 + \frac{p''(x)(x-u)}{p'(x)}$.[12]

We can then derive a second order differential equation, (A.1) in the Appendix, which

[12]Notice that $f_{x|z}(x(u)) = f_{u|z}(u(x))\left|\frac{\partial u(x,z)}{\partial x}\right| = f_{u|z}(u(x))\frac{\partial u(x,z)}{\partial x}$ since the SOC implies that $\frac{\partial u}{\partial x} \geq 0$ in interior optimum.

determines the optimal equilibrium responses $p(x|z)$ and $x(u|z)$ in audit group z using the expressions for $\mathrm{E}(y|x,z)$, $u(x|z)$, $\frac{\partial u}{\partial x}$, and the tax agency's first order condition. However, as some taxpayers pool at the lowest report, to obtain sufficient conditions for equilibrium, we must check the tax agency's first order condition at $x = \underline{u}$ separately as

$$\mathrm{E}(u|x = \underline{u}, z) = \frac{f^h_{u|z}(x)\, x + \int_{\underline{u}}^{u^{\text{pool}}} u \cdot f^d_{u|z}(u)\, \mathrm{d}u}{f^h_{u|z}(x) + \int_{\underline{u}}^{u^{\text{pool}}} f^d_{u|z}(u)\, \mathrm{d}u} = \frac{\lambda(z)\, c}{t + \theta t} + \underline{u}, \tag{5}$$

where u^{pool} is the residual income at which taxpayers (in this audit group) begin to pool at the lowest possible report.

Thus, given the equilibrium $\lambda(z)$, we can characterize the unique within-group equilibrium from Equations (5) and (A.1). By Equation (2), each $\lambda(z)$ corresponds to a required budget allocation, $B(z)$. Finally, the budget allocation across different z is pinned down by the requirement that the shadow value of increasing the budget, $\lambda(z)$, must be the same for all z, i.e., $\lambda(z) = \boldsymbol{\lambda}$, $\forall z$, for an interior solution. The shadow value, $\boldsymbol{\lambda}$, is pinned down by the requirement that the tax agency's overall budget, $\mathbf{B}$, may not be exceeded.

As mentioned above, the model contains Erard and Feinstein (1994) as a special case when attention is limited to a single audit group in which taxpayers without third-party income reports and the ratio of honest to dishonest taxpayers is constant on the domain of u. To illustrate, Figure 3 depicts the equilibrium for B at 10 percent, $\log(u) \sim \mathcal{N}(3.42, 0.3^2)$ truncated on $[20, 44]$, $Q = 0.4$, and $t = 0.5$.

Figure 3(a) shows the audit schedule, $p(x)$: it starts in $\underline{u}$, is downward sloping, and terminates in $p(\bar{x}) = 0$. This form balances the need to audit in order to raise revenue with the cost of doing so. The negative slope reflects the need to discourage high-income taxpayers from reporting too low incomes.

Figure 3(b) shows the amount of evasion as a function of true income. The linear increase in the first part of the graph reflects pooling of dishonest taxpayers: for a given audit schedule, there will be some level of residual income, u^{pool} in $[\underline{u}, \overline{u}]$, for which the most profitable report is $\underline{u}$. Consequently, all taxpayers with residual incomes $u < u^{\text{pool}}$

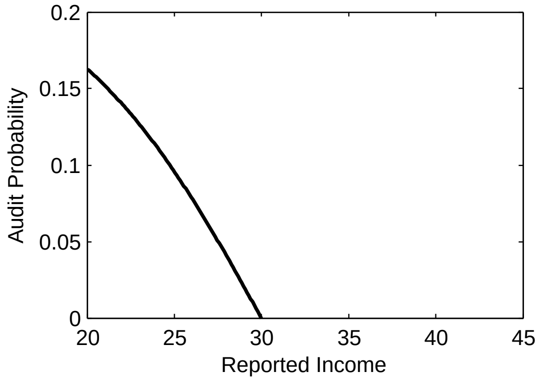

(a) The Optimal Audit Schedule, $p(x)$.

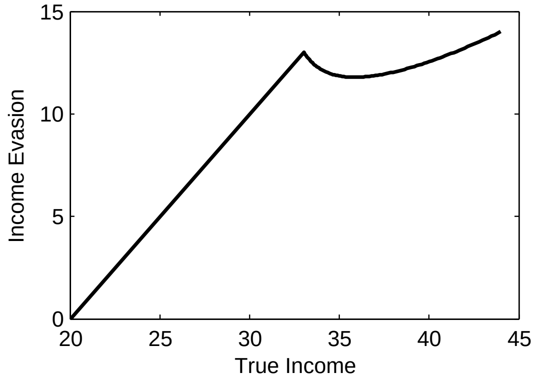

(b) Evaded Income, $u - x$, by True Income, y, for Dishonest Taxpayers.

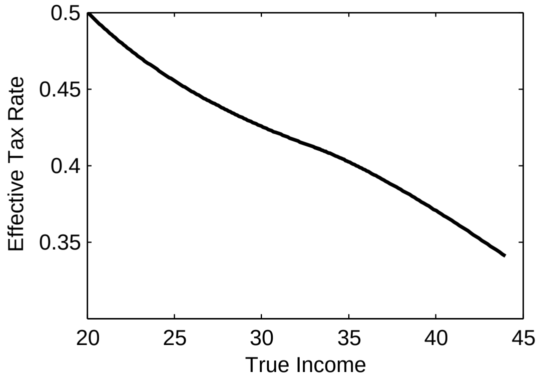

(c) Regressive Bias, $\tau - \tau^{\text{eff}}$, for Dishonest Taxpayers.

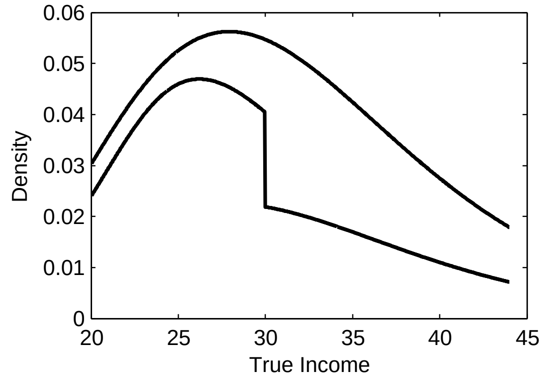

(d) Induced Reporting Behaviour. The lower curve graphs the density of reports by dishonest taxpayers, excluding the mass point at $x = \underline{u}$, while the upper curve graphs the true income distribution.

Figure 3. Equilibrium Responses and Tax Bias.
Notes: All panels display an example of equilibrium functions from the Erard and Feinstein (1994) model without third-party reporting. Equivalently, this could be an example of the solution for a particular z in our model including third-party reporting. This example is produced assuming $B = 10$ percent, $\log(u) \sim \mathcal{N}\left(3.42, 0.3^2\right)$ truncated on $[20, 44]$, $Q = 0.4$, and $t = 0.5$.

also report $x = \underline{u}$. Therefore, there will be a point mass in the induced distribution of reports, $f_x(x)$. After this pooling point, evasion falls rapidly in income until evasion again becomes increasing in income as the probability of detection becomes sufficiently low.

Figure 3(c) shows the effect of the optimal audit schedule on the *ex ante* effective tax rate, τ^{eff}, which is calculated as the ratio of expected payments (taxes and penalties) to true income

$$\tau^{\text{eff}} = \frac{p(x) \cdot (ty + \theta t (y - \tilde{y})) + (1 - p(x)) \cdot t\tilde{y}}{y}. \tag{6}$$

The declining profile of $p(x)$ together with the high propensity to evade taxes of high income taxpayers result in a negative relationship between the effective tax rate and income. Therefore, high-income taxpayers pay significantly less than the statutory tax rate, which, in the case of Figure 3(c), is $t = 0.5$, and we get regressively biased effective average tax rates.

Figure 3(d) shows the induced distribution of incomes and reports. The top graph is the original income distribution, which in this case is lognormal. The lower graph shows the distribution of induced reports, i.e., the equilibrium response of all taxpayers to the audit schedule. The right part of the graph is just a scaling of the original income distribution by Q while the left part is a weighted average of reports by honest and dishonest taxpayers. The whole graph is somewhat lower than the original income distribution as there is a mass point of dishonest taxpayers reporting at $\underline{u}$, the mass point being equal to the area between the graphs.

3 Data

SKAT's tax collection efforts extensively employ information reports by third parties. During some year t, incomes are earned and by the end of January in year $t + 1$, SKAT receives information reports from employers, banks, pension funds, and other entities, so-called third-party income reports. In general, all income received as salary, private/public pensions, honorarium, unemployment benefits, etc. is subject to third-party reporting as well as, e.g., mortgage interest payments and some capital income.[13] Self-employment income is rarely covered by information reporting except in cases where, e.g., remuneration is paid by a public institution. Third parties do not have discretion as to whether or not to supply SKAT with this information. The informational requirement is entirely related to the type of income.

By mid-March, SKAT sends out pre-populated tax returns based on third-party information and other information that they possess about the taxpayers, such as their residence and workplace for calculating commuting allowances. Subsequently, taxpayers have until May 1 to correct their tax return; in case of no corrections, the pre-populated tax return counts as final. After the deadline, SKAT's computerized system processes tax returns and attaches audit flags to returns that the system finds likely to contain errors. The system does not as such assign a probability of audit or rank tax returns according to their likelihood of containing errors but assigns a recommended action, i.e., "audit" or "do not audit". Briefly, the audit flag system relies on third-party income reports and also a collection of auditing "best practices" that could be converted to algorithmic form, e.g., specific tax return compositions indicative of misreporting, cut-off rules based on expected incomes conditional on third-party reported income, etc. The flag system consists of a large number of flags, each of which is intended to signal the likelihood of tax evasion on particular line-items or combinations of line-items. Although, the flag system operates for both wage earners and the self-employed, in practice, it is only used for wage earners as the predictive power of the audit flags for self-employed has been judged too low by SKAT. For the self-employed, further information is gathered on a case-by-case basis. Predominantly, SKAT uses correlates of true income such as bank deposits, consumption of housing, cars, and other durables to signal of the likelihood non-reported income. They may also seek information exchange with known tax shelters about foreign deposits or uncover such deposits indirectly by tracking purchases with foreign credit cards, but such information is much harder to gather. All in all, the workings of the audit regime is very different for the self-employed and much more resource intensive.

[13] Dividends are reported by third parties, whereas capital gains were not reported in 2006/2007.

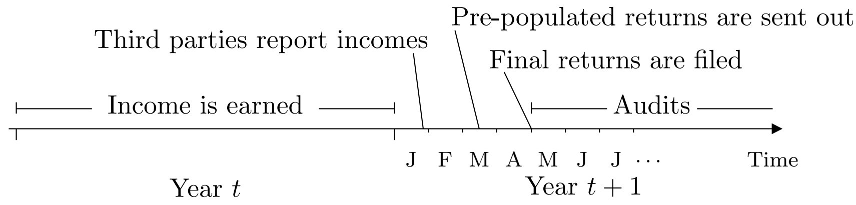

Figure 4. Tax Collection in Denmark – The Timing of Events.

After the tax returns have been processed, tax examiners assess the flagged returns and decide whether or not to initiate an audit based on the information available, local knowledge, and auditing resources. For wage earners, the information available is processed via the flag system and for the self-employed external information is gathered on an ad hoc basis. The process is depicted in Figure 4.

If an audit discovers underreporting, the taxpayer may pay the taxes owed immediately or postpone the payment at an interest. If the tax examiner views the underreporting as deliberate, the tax agency may impose a fine according to a fining scheme depending on the assessed intentionality of the misreporting.

3.1 Experimental Design

The data originates from an experiment conducted by SKAT in the years 2006–2008, originally analyzed in Kleven et al. (2011), and is in many ways comparable to the US Taxpayer Compliance Measurement Program. The experiment involved a stratified random sample of 17,764 self-employed individuals and 25,020 wage earners and recipients of public transfers in Denmark. In the present study, we use a sample of non-treated wage earners and recipients of public transfers (referred to as "wage earners") and a sample of non-treated self-employed for the fiscal year 2006.[14] The sample of wage earners is a stratified random sample of 10,740 Danish taxpayers, and the sample of the self-employed is a random sample (non-stratified) of 8,890 taxpayers.[15] The full populations of wage

[14]In the original study in Kleven et al. (2011), some taxpayers were subject to treatments. These taxpayers received notifications prior to filing their final tax returns, indicating that they would be audited with either 50 or 100 percent probability.

[15]Note the randomness of our sample as opposed to tax compliance data obtained from the regular audits that is heavily biased by over-sampling taxpayers who are likely to have misreported their income in either direction. The sampling strategy for wage earners involved a stratification on tax return complexity. For the self-employed no stratification scheme was employed.

earners and self-employed, respectively, where approximately 4.2 million and 400,000 in 2006. For each taxpayer, SKAT conducted an unannounced audit after the deadline for changing the tax return (May 1, 2007). The tax audits were comprehensive in the sense that SKAT examined all items on the tax return, demanding documentation for all items on which SKAT did not possess information. Moreover, SKAT made a significant effort to have tax examiners perform homogeneous audits by, e.g., organizing training workshops and distributing detailed audit manuals. The audits took up 21 percent of the resources devoted to tax audits in 2007.

Of course, it is unlikely that tax examiners find all hidden income, such as that stemming from cash-only businesses and other black market activities. We focus our attention on the detectable part of tax evasion given the methods available to SKAT and thus denote our empirical counterpart of true income "detectable income". In what follows, we will write true income when in fact we mean detectable income. In Section 4.3.4 we discuss the implications of this for our results.

For each taxpayer, we have income and tax records as reported by third parties, the final return as potentially changed by the taxpayer, and the post-audit return. In addition, the data contains information on the generated audit flags that would normally constitute a basis for selecting taxpayers for audits as well as a "compliance rating" reflecting the auditor's assessment of the degree to which discovered misreporting reflected deliberate fraud or accidental under/over-reporting.

3.2 The Tax System and Tax Compliance in Denmark

The Danish income tax system (in 2006) operates with many different measures of income. Here, we will provide the headlines; see Table 1 for details. Labor market income, i.e., salary, fringe benefits and other earned income, are taxed proportionally by a labor market tax of 8 percent while an earned income tax credit (EITC) of 2.5 percent is provided for labor market income up to 292,000 DKK.[16] Capital income is a net concept, and different tax rates apply depending on whether net capital income is positive or negative. For most

[16] Approx. 49,000 USD (1 USD $\approx$ 6 DKK in 2006).

Table 1. An Overview of the Danish Tax System, 2006.

Tax	Tax base	Bracket (DKK)[a]	Rate (pct.)
Labor market tax	Labor inc.	none	8.0
EITC	Labor inc.	up to 292,000	2.5
Bottom tax	Personal inc.+ max(cap.inc., 0)	38,500–	5.5
Middle tax	— // —	265,500–	6.0
Top tax	— // —	318,700–	15.0[b]
Local taxes	Taxable inc. (=pers.inc.+cap.inc.−deductions)	38,500–	33.3[c]
Stock income tax	Stock inc.	0–44,300; 44,300–	28.0; 43.0

[a]1 USD ≈ 6 DKK (in 2006).
[b]The top tax rate may be lowered by the "tax ceiling" that limits the sum of state taxes (bottom, middle and top) and local taxes (excl. church taxes) to 59 percent. In the average municipality the tax ceiling lowers the top rate by 0.08 percentage points.
[c]In the avg. municipality and county incl. optional church tax of on avg. 0.74.

taxpayers, net capital income is negative due to interest payments on mortgages. Central government taxes (bottom, middle and top tax) are levied on the so-called "personal income", which, in addition to positive net capital income, consists of labor market income plus social transfers and pensions less labor market taxes and some pension contributions. Central government taxes constitute a progressive tax scheme with a personal allowance and three brackets. Local taxes (county and municipality) are levied on "taxable income", which is similar to the central government tax base except that it allows for negative net capital income deductions and other deductions such as transport allowances. In this way, Denmark has a version of the Nordic dual income tax;[17] negative capital income is taxed at a flat rate, whereas positive capital income is taxed progressively just as regular income. Stock income (dividends and capital gains) is subject to a two-rate scheme with the high rate setting in at 44,300 DKK.

Table 2 presents some descriptive statistics on major income components for the two samples of wage earners and self-employed, respectively. The table shows sample means with standard errors of means in parentheses – all numbers for wage earners are calculated accounting for the stratification scheme. Column (1) presents pre-audit figures measured at the deadline, May 1, and column (5) shows figures reported by third-parties. Self-reported figures (the difference between (1) and (5)) are shown in column (6). Negative figures mean that taxpayers on average adjust the number downwards to less than what

[17]For a discussion of the Nordic dual income tax., see e.g. Nielsen and Sørensen (1997).

third-parties have reported. Columns (2)–(4) describe how the figures in (1) were adjusted by the tax examiners during the audits. Columns (3) and (4) split the audit adjustments into positive (meaning underreporting) and negative (meaning overreporting) adjustments while column (2) holds the average net adjustment, i.e., the sum of (3) and (4).

Panel A of Table 2 shows figures on total income and total taxes for wage earners. The former is defined as the sum of personal income, capital income, stock income, self-employment income, and foreign income less deductions. Pre-audit total income is on average a little less than 200,000 DKK with a significantly positive net adjustment from SKAT of almost 1,700 DKK. The positive net adjustment reflects an asymmetry in the reporting behavior with underreporting being more than ten times as high as the overreporting on average. Third-party reported total income is slightly higher than pre-audit total income mainly due to deductions not included in the third-party reports, implying a negative residual (i.e., self-reported) total income.

Panel B features a decomposition into main income components for wage earners. The asymmetry in the over- and underreporting found for total income is noticeable for all components.[18] Not surprisingly, the greatest relative amount of underreporting is found on items least subject to information reporting. Self-employment income tops the list with underreporting amounting to 18.5 percent of the mean post-audit self-employment income level followed by stock income (6.8 percent), deductions (2.3 percent), and the rest being less than 2 percent.

In Panel C we show descriptive statistics for the sample of self-employed taxpayers. As a decomposition into income components has not been possible, we only show numbers for total income and total taxes.[19] As with wage earners, we find a pronounced asymmetry in net audit adjustments corresponding to much higher underreporting compared to overreporting for the self-employed. The main difference compared to wage earners is

[18]Foreign income is the exception. Here, the average net adjustment is negative, corresponding to overreporting on average, yet, the adjustment is not significantly different from zero. The likely reason is that there are few cases of foreign income, and the variation in adjustments performed by SKAT is dominated by correction of mistakes.

[19]During the experiment, tax corrections concerning the interplay of the business and private side for the self-employed where not included in the data. This was remedied at the aggregate level for total income and total taxes but not for separate income components. Therefore, for the self-employed we only include descriptive statistics for aggregate income and tax measures.

Table 2. Tax Compliance in Denmark, Income Year 2006.

	Reported income	Net audit adjustment	Under-reporting	Over-reporting	Third-party rep. inc.	Self-rep. inc.
	(1)	(2)	(3)	(4)	(5)	(6)
A.			Wage earners			
Total Income	193,277	1,664	1,825	-161	195,618	-2,341
	(1,906)	(480)	(479)	(22)	(1,844)	(584)
Total Tax	63,178	636	695	-59		
	(841)	(246)	(246)	(9)		
B.			Income components			
Earnings	156,127	672	683	-11	155,987	140
	(2,275)	(203)	(203)	(6)	(2,217)	(559)
Personal inc.	209,232	1,137	1,195	-58	209,726	-494
	(1,950)	(480)	(479)	(17)	(1,886)	(573)
Capital inc.	-10,884	142	198	-56	-11,308	424
	(272)	(27)	(24)	(11)	(266)	(81)
Deductions	-9,264	143	213	-70	-5,605	-3,659
	(178)	(28)	(26)	(11)	(85)	(144)
Stock inc.	3,612	239	262	-24	2,797	815
	(546)	(40)	(39)	(10)	(502)	(188)
Self-empl. inc.	103	21	23	-2	8	95
	(60)	(8)	(8)	(1)	(4)	(60)
Foreign inc.	479	-18	6	-25	0	479
	(92)	(19)	(4)	(19)	.	(92)
C.			Self-employed			
Total Income	298,388	21,480	22,697	-1,217	157,285	141,103
	(8,321)	(1,912)	(1,905)	(145)	(6,445)	(5,534)
Total Tax	124,392	8,719	9,089	-371		
	(4,423)	(609)	(606)	(50)		
D.			Wage earners and self-employed			
Total Income	202,310	3,367	3,619	-252	192,324	9,987
	(1,883)	(469)	(467)	(24)	(1,774)	(715)
Total Tax	68,439	1,331	1,416	-86		
	(858)	(231)	(231)	(9)		

Notes: Panels A and B show descriptive statistics for a stratified random sample of 10,740 taxpayers denoted as wage earners (incl. unemployed, pensioners, etc.). Due to the stratification strategy employed by SKAT, the sample contains 74.6 percent "heavy" taxpayers (i.e., with high-complexity tax returns) and 25.4 percent "light" taxpayers, whereas the population has 32.6 percent heavy taxpayers and 67.4 percent light taxpayers. In Panel C the sample consists of 8,890 randomly selected self-employed taxpayers. No stratification was employed. Panel D provides descriptive statistics for wage earners and self-employed combined using population weights.
Total income is defined as personal income + capital income − deductions + stock income + self-employment income + foreign income. The decomposition in Panel B is only available for the sample of wage earners. In the table, deductions are given as a negative amount. Reported income is the sum of third-party reported income and self-reported income. Standard errors of means in parentheses. All estimates for wage earners are population weighted.
All amounts in DKK (1 USD $\approx$ 6 DKK in 2006).

spelled out in the average level of self-reported income. Income sources of self-employed are to a much lesser extent covered by the system of third-party reporting, resulting in an almost even split between income reported by third parties and self-reported income. This provides SKAT with a much greater challenge in discovering unreported income.

We get a further idea as to where the opportunities to evade taxes are prevalent by looking at taxpayers' behavior and conditioning on the informational environment. In Table 3 we separate taxpayers according to whether or not their entire income was reported to the tax agency by a third party. Panel A shows the shares of under-/overreporting and correct reports for each sample (wage earners and self-employed, respectively). All figures in the table are calculated accounting for stratification whenever applicable. The overall population weighted share of compliers, given by wage earners not underreporting, amounts to approximately 94 percent for wage earners. For the self-employed, approximately 65 percent comply. To address taxpayers with ample opportunity to evade taxes, Panel B shows shares of particular groups conditional on whether or not their entire income is reported by a third-party (standard errors in parentheses). For example, less than 2 percent of wage earners with all income reported by third parties underreport taxes. For wage earners with some income not reported by third parties, this share is much higher, yet a substantial share of over 80 percent (depending on the definition of compliance) are found to comply with the tax laws despite having ample opportunity to evade.

Only few self-employed taxpayers (3.4 percent) have their entire income reported by third parties, underlining the tax agency's challenge in securing tax revenue from these taxpayers. Further, almost 35 percent are found to underreport their taxes. The share of self-employed who do not underreport their taxes is again high (95 percent) for those with all income reported by third-parties and much lower (64 percent) for those with some income not covered by the system of third-party reporting, albeit still a substantial share comply with given tax laws. Strikingly, wage earners and self-employed who have all their income reported by third parties do not differ significantly in reporting behavior, whereas this is not the case when some income is not reported by third parties.

Table 3. Reporting Behavior of Danish Wage Earners and the Self-Employed, 2006 Incomes.

	Wage earners		Self-employed	
Observations	10,740		8,890	
Entire income reported by third-parties?	Yes	No	Yes	No
A.	Share	Share	Share	Share
# underreported	0.010	0.049	0.002	0.346
# correct	0.653	0.269	0.032	0.570
# overreported	0.003	0.016	0.000	0.048
Total reports	0.665	0.335	0.034	0.966
B.	Share of sub-sample	Share of sub-sample	Share of sub-sample	Share of sub-sample
Correct reports	0.979	0.809	0.943	0.590
	(0.002)	(0.011)	(0.055)	(0.006)
Not underreporting	0.984	0.855	0.950	0.640
	(0.002)	(0.010)	(0.055)	(0.007)
"Honest" taxpayers*	0.988	0.901	0.957	0.690
	(0.002)	(0.008)	(0.055)	(0.007)

Notes: Standard errors of fractions in parentheses. The sample of wage earners is a stratified random sample. Fractions and standard errors are calculated subject to the stratification scheme. "Wage earners" also include recipients of benefits. The sample of self-employed is a non-stratified random sample.
*Calculated imposing the assumption that unintentional underreporting is as frequent as (unintentional) overreporting – i.e., symmetry in reporting errors. For example, for the self-employed (right-most column), the (unstratified) calculation is simply $(0.570 + 2 \cdot 0.048)/0.966 \approx 0.690$. For wage earners, we provide a population weighted estimate.

3.3 Effective Tax Rates

To address the distortion of tax evasion/enforcement of the effective tax system, we need a measure of *ex ante* effective average tax rates calculated directly from data. As with Equation (6), an appropriate measure of effective average tax rates must take into account the extent of tax evasion, the risk of detection, and the penalties paid in case of detection, all of which affect expected payments to the tax agency. As such, we must restrict our analysis of this phenomenon to the sample of wage earners for whom we have information about the enforcement regime. The *ex ante* effective average tax rate can be calculated from data as

$$\tau^{\text{eff}} = \frac{f \cdot \left(T + \Theta\left(T - \tilde{T}, I\right)\right) + (1 - f) \cdot \tilde{T}}{Y}, \tag{7}$$

where f is the probability of getting caught, T and $\tilde{T}$ are taxes on true and reported income, respectively, Y is true income, and $\Theta(\cdot,\cdot)$ is a nonlinear function describing the penalty for underreporting taxes as a function of underreported taxes and the assessed intentionality of evasion, I. With a probability f, evasion is detected and the taxpayer pays the full taxes due plus a penalty that is proportional to the amount of taxes evaded. With a probability $1-f$, evasion goes undetected and the taxpayer only pays taxes on reported income. As tax evasion diminishes or as the risk of detection or the penalty increase, the effective average tax rate will increase, *ceteris paribus*.

We denote by τ the nominal average tax rate, defined in the usual way, $\tau = T/Y$. As a matter of convenience, we define the tax rate bias as $\tau - \tau^{\text{eff}}$. This allows us to compare how much statutory and effective tax rates differ when both vary across individuals in the sample. Although we focus on individuals for which the data reveals some underreporting, (7) is equally valid for taxpayers not engaging in tax evasion. For these individuals, taxes due on reported and actual income are the same, and the effective average tax rate is simply the average tax rate, $\tau^{\text{eff}} = \tau$.

Y, T, and $\tilde{T}$ are observed in the data as post-audit total income and taxes, and pre-audit taxes.[20] We use SKAT's audit flag system as a proxy for the probability of getting caught for wage earners. Not all taxpayers with flags are audited, so we assume that the probability is proportional to the number of flags assigned to a tax return.[21] Specifically, we calculate our proxy for the probability of detection simply as the ratio of flags assigned to a tax return to the maximal number of flags assigned to any tax return. With this approach, the audit rate among wage earners is 3.3 percent. This is slightly lower than the total population audit rate of 4.2 percent reported by Kleven et al. (2011). As this rate includes audits of the self-employed, who, presumably, are audited relatively more intensively, the average audit rate suggested by our proxy seems more or less reasonable.

We specify the penalty function, $\Theta\left(T-\tilde{T}, I\right)$, using the actual rules for calculating penalties for tax evasion and the compliance rating system applied by the tax examiners

[20]Recall the definition of total income as the sum of personal income, capital income, stock income, self-employment income, and foreign income less deductions.

[21]Alternatively, this can also be interpreted as an assumption that each part of the tax return, to which an audit flag corresponds, is audited with probability 1.

during the audits. In Denmark, evasion penalties are calculated as a factor on taxes evaded; that factor, however, varies for the amount evaded and the intentionality of evasion as assessed by the auditor. In the case of intentional tax evasion, the fine is calculated as 1 times evaded taxes under 30,000 DKK and 2 times the evaded taxes exceeding 30,000 DKK. In the case of gross negligence, the rates are instead 0.5 times evaded taxes not exceeding 30,000 DKK and 1 times evaded taxes exceeding 30,000 DKK. Fortunately, the compliance ratings in the data are exactly intended to measure the degree of intentionality of uncovered tax evasion. Compliance ratings take on values in $\{0, 1, 2, \ldots, 6\}$ indicating decreasing degrees of intentionality of misreporting. According to this classification, compliance ratings of 0, 1, or 2 signify deliberate tax evasion, whereas $3, \ldots, 6$ signify gross negligence (approaching 3) or innocent mistakes (approaching 6). Using these classifications, we can accurately calculate the penalty rate applicable for each individual tax evader.[22]

4 Calibration and Results

Due to the considerable detail of our data, we can construct a set of parameters for the purpose of simulating the model that are internally consistent, i.e., they all derive from the same data set. Using the samples of wage earners and the self-employed we approximate penalty and tax rates from the actual tax system. The parametrized share of honest taxpayers we allow to differ between the groups of wage earners and the self-employed to account, in some measure, for self-selection into these employment categories. For the same reason, we also estimate the bivariate income distributions separately for the two groups. As we calculate below, the share of honest taxpayers is indeed much lower for the self-employed corresponding to the intuition that some people may self-select to exploit more ample evasion opportunities. Finally, for each group, we calibrate the audit budget to match simulated average tax evasion among evaders to observed average evasion. Without loss of generality, we can normalize the per-audit cost, c, to 1 such that

[22] Assuming, e.g., that innocent mistakes (rated 6) are not penalized or that the threshold in compliance ratings between intentional evasion and gross negligence is between 1 and 2 or between 3 and 4, turns out not to affect the results we present in Section 4.

overall budget parameters **B** can be interpreted as the share of the population subject to audit within the groups of wage earners and the self-employed.

4.1 Calibration

4.1.1 Income Distributions

We use the taxpayer data to construct the income distributions needed in the model. As income measure we use total income defined as the sum of personal income, capital income, stock income, self-employment income, and foreign income less deductions.

In principle, the densities of honest and dishonest taxpayers can be estimated separately but with the size of our data set this would introduce a large element of uncertainty in estimates of dishonest taxpayers. Instead, we follow Erard and Feinstein (1994) closely and estimate a common distribution for both honest and dishonest taxpayers with the only difference being a mass point of honest wage earners for whom true income is perfectly predicted by third-party reported income, i.e., $u = 0$, which allows this mass point to vary in z. This is important because richer wage earners are much more likely to have non-zero residual income than poorer wage earners. However, for the self-employed there are very few individuals without some residual income and we can estimate income distributions without accounting for a mass point. In practice, to fit the simultaneous distribution of z and u, we exclude any honest taxpayers in $u = 0$ and fit a mixed lognormal distribution.[23] The distribution of the mass point of wage earners at $u = 0$ across z is estimated separately.

The exact characteristics of this distribution is documented in the Appendix. Briefly, the variance of $u|z$ is generally increasing in z; however, the taxpayers with very low or negative z have relatively complicated income compositions resulting in high variance of $u|z$ and, for wage earners, a relatively small mass point at $u = 0$.

[23]Our results do not appear to alter significantly if, instead, a kernel estimation is used. However, kernel densities are inconvenient as they allow for "troughs" of zero density in the interior of $[\underline{u}, \overline{u}]$ which may cause our algorithm to fail. By using a sufficient number of component distributions in the mixed lognormal distribution, the difference between this distribution and a bivariate kernel distribution becomes negligible.

4.1.2 Honesty

With our simplified version of conditional densities, we can write $f^h_{u|z} = Qf_{u|z}(u) + \mathbf{1}_{(u=0)}M(z)$ and $f^d_{u|z} = (1-Q)f_{u|z}(u)$, where $\mathbf{1}_{(\cdot)}$ is the indicator function. Thus, for $u \neq 0$ the share of honest taxpayers is Q, whereas for $u = 0$ it is $Q + M(z)$, where $M(z) \geq 0$ is the mass point at $u = 0$ for some level of third-party reporting, z. To determine an appropriate value of the parameter, Q, we must account for the fact that, in reality, some taxpayers seem to make reporting mistakes. For example, in the data some reports are adjusted downward by the auditor, which means that, in the absense of an audit, the taxpayer would have payed more than intended by the statutory tax system.

We approach the problem in the following way. First, we assume that no taxpayer will try to evade taxes on income that is reported by a third party (this assumption is bourne out in the data for wage earners as shown in Table 3). Secondly, in keeping with the model, we disregard the fact that some taxpayers make reporting mistakes. A revenue maximizing tax agency is indifferent about the motivation for underreporting and about overreporting.[24] As a consequence, taxpayers reporting too large taxable incomes are treated as if they are exactly compliant and taxpayers that underreport taxable incomes by mistake are treated as tax evaders. Then we separate taxpayers by whether they underreported taxes (non-compliant taxpayers, $x < u$) or reported correctly/overreported taxes (compliant taxpayers, $x \geq u$). Compliant taxpayers are then decomposed into those with zero residual income and non-zero residual income. We define the parameter Q as the ratio of compliant taxpayers with non-zero residual income to the total number of taxpayers with non-zero residual income in the sample. The idea is that having some income not subject to third-party reporting provides taxpayers with ample opportunity for evasion. By not seizing the opportunity, they reveal themselves as being honest in the present context. Table 3 shows this decomposition. First, note that among wage earners whose entire income is reported by third parties, the compliance rate is 97.9 percent. Among those wage earners that have some of their income not reported by

[24] We do not consider the, rather implausible, scenario that the tax agency might refrain from auditing certain groups because this would reveal overreporting by some taxpayers thus lowering collected revenue.

third parties, the compliance rate is 80.9 percent. The number of honest taxpayers is the sum of those reporting correctly and those overreporting by mistake, which corresponds to $Q = 85.5$ percent.[25] The residual consists of both dishonest wage earners and wage earners underreporting by mistake whom we cannot distinguish. To partially control for self-selection into occupations according to a taxpayer's proclivity to evade taxes, we calculate Q separately for the self-employed as shown in Table 3. The resulting value, $Q = 64.0$ percent, is indeed substantially lower and suggests that this distinction is important.

4.1.3 Penalty

The model has a fixed penalty factor, θ, as opposed to the more complicated penalty function, $\Theta(\cdot,\cdot)$, from Section 3. We approximate an appropriate value of θ by calculating the average penalty rate for the sample of tax evaders accounting for stratification between light and heavy taxpayers within the group of wage earners and for the relative shares of wage earners and self-employed in the population. We take a simple approach and use the OLS slope coefficient between calculated penalties, $\Theta(\cdot,\cdot)$, and underreported taxes as our value of θ. The resulting penalty rate on underreported taxes is 1.15.

4.1.4 Tax Rates

We estimate a marginal tax function, $t(z)$, using local mean smoothing of marginal tax rates on the entire sample of wage earners and self-employed accounting for stratification of light and heavy taxpayers in the group of wage earners and the relative shares of wage earners and self-employed in the population. We allow the approximated tax rates to vary in z to partially account for the progressiveness of Danish income taxes. Because our data set contains all line items, we can calculate each taxpayer's marginal tax rate on all income components, such as earnings, capital income, stock income, etc. For each component, we calculate marginal taxes with respect to reported income. To obtain an average marginal effective tax rate, conditional on z, we then weight marginal taxes of

[25] Of course, we also account the for the sample stratification in calculating Q.

different components according their relative prominence on a taxpayers tax return (i.e., before the taxpayer is audited).

4.2 Simulation Strategy

An individual solution, $\left(p, \frac{\partial p}{\partial x}\right)$, to Equation (A.1) in the Appendix that corresponds to a particular z is found numerically using methods of Ordinary Differential Equations (ODE). The solver is initialized using $p(\bar{x}) = 0$ and $p'(\bar{x}) = \left(\frac{1}{1+\theta}\right) / (\bar{u} - \bar{x})$, where $\bar{x} \equiv x(\bar{u})$. Thus, starting at the end-point of the equilibrium-path audit probabilities, a numerical solver finds values in steps until $\underline{u}$ is reached, ensuring that the taxpayers' as well as the tax agency's optimality conditions are met for reports $x \in (\underline{u}, \bar{x}]$. However, since a positive mass of taxpayers are pooling their reports at $x = \underline{u}$, the expectation $\mathrm{E}(u|x, z)$ is not differentiable in this point. Therefore, we check that the tax agency's FOC is met in the pooling point separately after finding some candidate solution, cf. (5).

The difficulty in identifying equilibria in this model stems from *a priori* indetermination of λ and $\bar{x}$: we must satisfy $\mathrm{E}(u|x = \underline{u}, z) - \underline{u} = \frac{\lambda c}{t+\theta t}$, which depends on both variables. Our solution method, the so-called shooting method for parametrized ODEs, searches the space of possible $(\lambda, \bar{x})$ for candidate solutions, for each checking that the tax agency's optimization constraints are satisfied on the entire domain of x, until satisfactory solutions are found. The optimal budget allocation, which in our simulations is always interior, equates shadow prices of increasing the budget density across levels of z.

While mathematically and intuitively z is naturally understood to be a continuous variable described by the simultaneous distribution of u and z, we approximate the optimal allocation of the total audit budget on the domain of z by constructing a representative, evenly spaced grid. We provide detailed documentation of the numerical implementation in the Appendix.

We have estimated $t(z)$, θ, Q, and the income distribution from data. Thus, the remaining free parameter is the budget value, $\mathbf{B}$, which we do not know. Since the mean level of evasion is inversely proportional to total tax revenue, it is monotonically declining in $\mathbf{B}$. To calibrate $\mathbf{B}$, we use the estimated income distribution to simulate a population of taxpayers: we vary **B** until the average level of evasion for tax evaders matches the level observed in the data, approximately DKK 8,312 for wage earners and DKK 25,991 for the self-employed. The resulting budget values are $\mathbf{B} = 0.0412$ and $\mathbf{B} = 0.4565$, respectively.

4.3 Results

As mentioned, we calibrate the model to the average level of evasion among tax evaders in the data. The match between data and simulations may seem trivial as it is imposed by the calibration procedure. However, in the context of the economic literature on tax evasion, being able to match a structural model to moments of the data for reasonable parameter values is, to our knowledge, novel. For example, Alm, McClelland, and Schulze (1992) argue that observed evasion is too low to be explained by a model of actual audit and penalty regimes. Our analysis lends support to the argument of Andreoni, Erard, and Feinstein (1998) and Slemrod (2007) that third-party reporting and tax-return-dependent audits can explain a substantial part of observed evasion. However, in accordance with Feld and Frey (2002), our analysis also requires us to take into account the substantial number of taxpayers that report honestly despite incentives to evade.

4.3.1 Tax Evasion and Enforcement

In Figure 5 we compare the observed distribution of flags across third-party reported income with the optimal distribution obtained in the simulations. In panel (a) we show, for each individual, third-party reported income (in '000 DKK) and the ratio of flags to the maximally observed number of flags assigned to any return. In addition, we show the local average ratios and 95 percent confidence bounds using local mean smoothing. In panel (b) we show individual and average observations of audit probabilities from simulated data. Generally, the audit intensity is increasing in third-party reported income, reflecting the fact that higher-income taxpayers find it relatively easier to evade taxes since the conditional variance of true residual income is larger. As we do not know how the number of flags assigned to tax returns translates into the likelihood of being audited, it

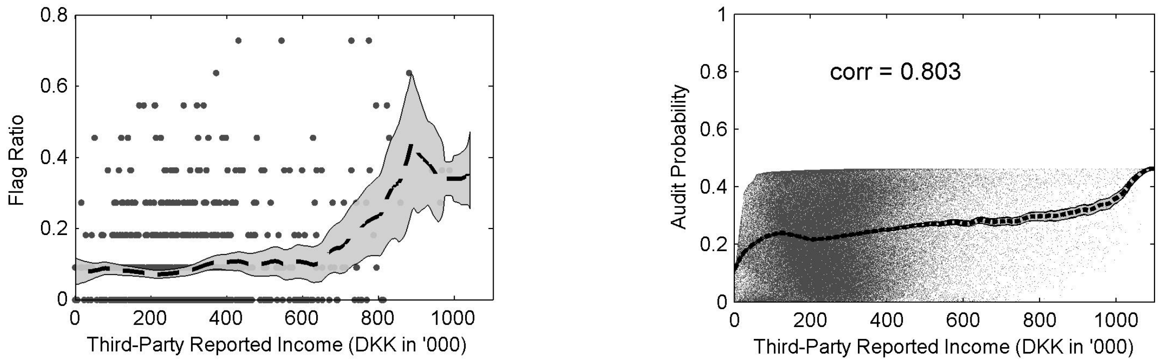

(a) Observed Ratio of Audit Flags per Taxpayer.

(b) Simulated Optimal Audit Probability.

Figure 5. Observed and Simulated Optimal Audit Intensity Across the Distribution of Third-Party Reported Income.

Notes: Panel (a) shows, for the subsample of only tax evading wage earners and recipients of benefits (905 obs.), the number of flags per taxpayer as a share of the maximally observed number of flags across the distribution of third-party reported income, z. Panel (b) shows the simulated audit probability ($\sim$194,000 obs.) across the distribution of third-party reported income, z. In both panels, the dotted lines give the local average of the observations together with 95 percent confidence bands using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. The local mean smoothing in Panel (a) does not account for the stratification scheme. The simulated data in Panel (b) is not stratified. Income is defined as the sum of all income less deductions and is measured in '000 DKK. 1 USD $\approx$ 6 DKK (in 2006). In Panel (b), the budget is allocated such that approximately 4.1 percent of all wage earners and recipients of benefits are audited.

is not surprising that there is a level difference between the two graphs. This reflects the fact that our minimal assumptions proxy for the empirical audit probability suggests an audit rate among wage earners of 3.3 percent, whereas the audit rate required to calibrate the model is 4.1 percent. Nonetheless, the graphs have very similar profiles. Both are increasing in third-party reported income and the audit intensity is especially high in the right tail of the distribution. This is bourne out in a correlation coefficient between local averages of 0.803.

Similarly, Figure 6 shows the empirical and simulated covariation of reported residual incomes (x), denoted in DKK in '000, and the probability of audit, which in panel (a) is proxied by the ratio of the number of flags assigned to a tax return to the maximally observed number of flags assigned to any tax return. As in Figure 5, there is a level difference between the two graphs. However, under our minimal assumption that the number of flags is positively correlated with the actual likelihood of audit, Figure 6(a) does suggest that the actual likelihood of an audit is distributed across the distribution of

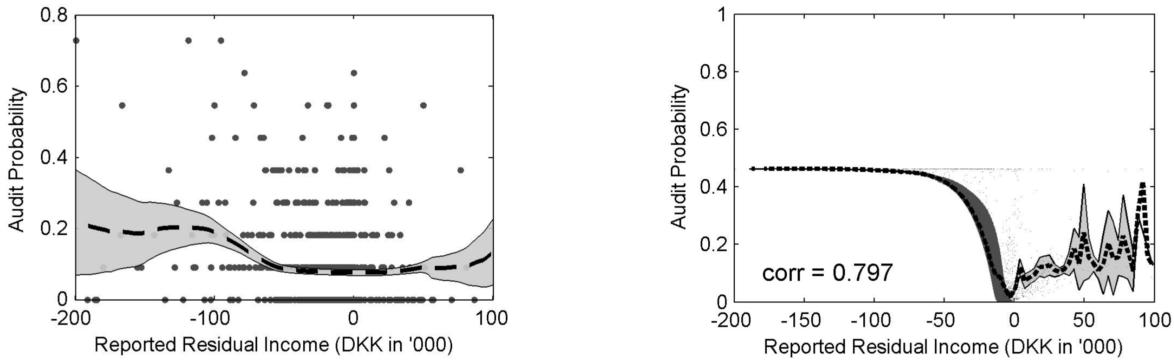

(a) Observed Ratio of Audit Flags per Taxpayer.

(b) Simulated Optimal Audit Probability.

Figure 6. Observed and Simulated Optimal Audit Intensity Across the Distribution of Reported Residual Income.
Notes: Panel (a) shows, for the subsample of only tax evading wage earners and recipients of benefits (905 obs.), the number of flags per taxpayer as a share of the maximally observed number of flags across the distribution of reported residual income, x. Panel (b) shows the simulated audit probability ($\sim$194,000 obs.) across the distribution of reported residual income, x. In both panels, the dotted lines give the local average of the observations together with 95 percent confidence bands using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. The local mean smoothing in Panel (a) does not account for the stratification scheme. The simulated data in Panel (b) is not stratified. Income is defined as the sum of all income less deductions and is measured in '000 DKK. 1 USD $\approx$ 6 DKK (in 2006). In Panel (b), the budget is allocated such that approximately 4.1 percent of all wage earners and recipients of benefits are audited.

reported residual income in a manner broadly consistent with a revenue maximizing tax agency.[26] This conclusion is reinforced by the relatively high correlation (0.797) between local averages of the share of flags and simulated audit probabilities in the distribution of reported residual incomes.

The simulations accurately reproduce the covariance structure of tax evasion with respect to the composition of the tax return in terms of third-party reported income and residual income. Figure 7 shows empirical and simulated covariation of tax evasion $(u - x)$ and residual income (u), denoted in DKK '000 for both wage earners and the self-employed. For each panel, we show individual data points and local averages and 95 percent confidence intervals across the domain of residual incomes using local mean smoothing. As shown in Panel (b) and (d), the local averages of simulated tax evasion

[26]Note that the increasing average probability of audit for $x > 0$ is perfectly consistent with audit probability functions being strictly descreasing, conditional on z. In the simulations, the average audit probability is increasing for $x > 0$ because the equilibrium audit intensity and the variance of $u|z$ are increasing in z. Therefore, the higher is a taxpayer's z, there more likely it is, on average, that he is audited which, in equilibrium, lessens the degree to which he evades taxes, making it more likely that he reports a positive residual income.

are highly correlated (correlation coefficients 0.951 and 0.966, respectively) with the local averages of observed tax evasion for both wage earners and the self-employed. Moreover, except for a slight clustering of wage earners with small negative residual incomes but relatively large degrees of evasion in the observed data, the distribution of individual data points also closely resembles that observed in the data. Although the self-employed evade more taxes on average, in neither data nor simulations do the self-employed appear to be more prone to evasion on the margin. Rather, the self-employed evade more taxes because they tend to have larger incomes and because less of that income is revealed by third parties. Finally, in equilibrium, tax evasion for the self-employed is curtailed to a large extent by intensive auditing.

In conjunction, Figures 5-7 suggest the direct evidence of tax evasion and the indirect evidence on the Danish tax agency's enforcement strategy is consistent with our theory of rational tax evaders and a revenue maximizing tax agency. Moreover, since, as shown in Figure 5, the budget intensity increases with third-party reported income, we should expect Schotchmer's conjecture of progressive/regressive bias conditional on residual/third-party reported income to be bourne out in both data and simulations.

4.3.2 Effective Tax Rate Bias

We calculate the bias of effective average tax rates as described in Section 3, $\tau - \tau^{\text{eff}}$, for data using the actual tax and penalty systems while for simulations using our approximations of a constant penalty rate, θ, and a set of constant marginal tax rates, t_z, that vary with third-party reported income.

In Figure 8 we display for each individual third-party reported income and our calcutation of effective tax rate bias. Panel (a) shows observations from the data set and Panel (b) shows simulated data. In each panel we also show local averages calculated using local mean smoothing. Both data and simulations exhibit effective average tax rates that are progressively biased with the bias decreasing towards 0 as third-party reported income increases. Moreover, the estimated local averages are highly correlated (correlation coefficient 0.974).

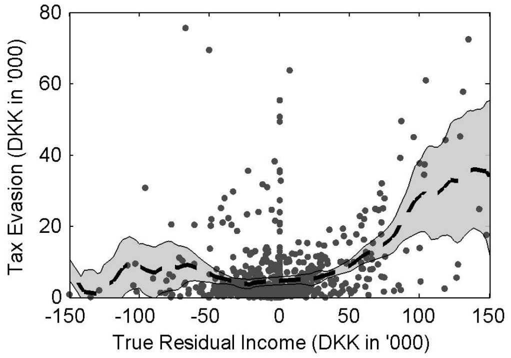

(a) Observed Tax Evasion, Wage Earners.

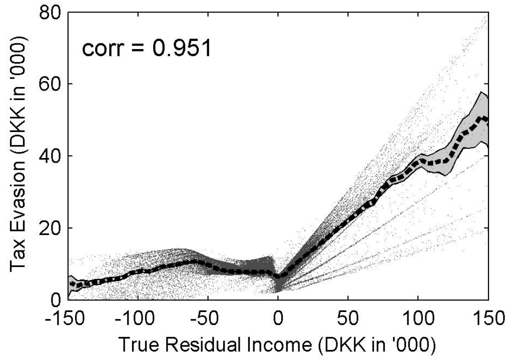

(b) Simulated Tax Evasion, Wage Earners.

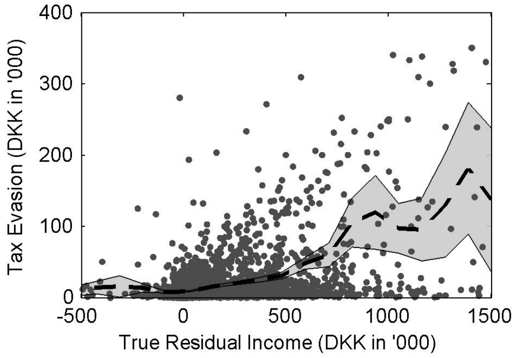

(c) Observed Tax Evasion, Self-Employed.

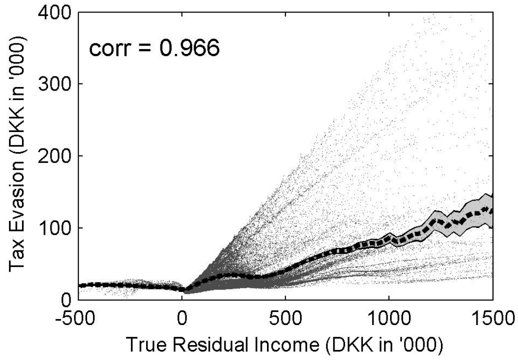

(d) Simulated Tax Evasion, Self-Employed.

Figure 7. Observed and Simulated Tax Evasion Across the Distribution of True Residual Income.

Notes: Panels (a) and (c) show observed tax evasion across the distribution of true, i.e., post-audit, residual income, u, for wage earners (905 obs.) and self-employed (2,980 obs.), respectively. Panels (b) and (d) show simulated tax evasion across the distribution of true, i.e., post-audit, residual income, u, for wage earners ($\sim$194,000 obs.) and self-employed ($\sim$190,000 obs.), respectively. In both panels, the dotted lines give the local average of the observations together with 95 percent confidence bands using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. The local mean smoothing in Panels (a) and (c) does not account for the stratification scheme. The simulated data in Panels (b) and (d) is not stratified. Income is defined as the sum of all income less deductions and is measured in '000 DKK. 1 USD $\approx$ 6 DKK (in 2006). In Panel (b), the budget is allocated such that approximately 4.1 percent of all wage earners and recipients of benefits are audited. In Panel (d), the fraction of self-employed taxpayers audited is approximately 45.7 percent. Note, however, that the self-employed make up a much smaller group (approx. 400,000) compared to wage earners (approx. 4.2 million).

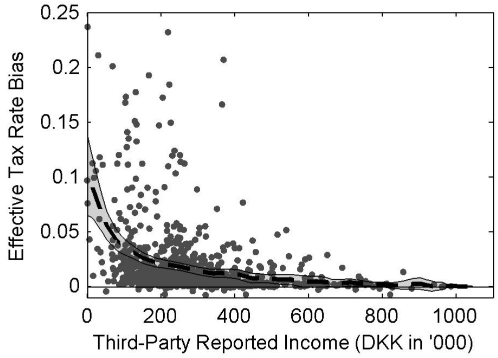

(a) Observed Tax Bias for Tax Evaders Across the Distribution of Third-Party Reported Income.

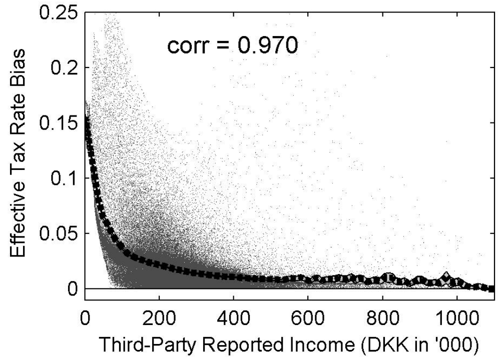

(b) Simulated Tax Bias for Tax Evaders Across the Distribution of Third-Party Reported Income.

Figure 8. Observed and Simulated Progressive Bias in Third-Party Reported Income.
Notes: The effective tax rate bias, $\tau - \tau^{\text{eff}}$, is the difference between the average statutory tax rate and the average effective tax rate as implied by the tax system, tax enforcement, and tax evasion behavior. Panel (a) shows the observed tax bias as a function of third-party reported income, z, for the subsample of tax evading wage earners and recipients of benefits (900 obs.). Tax rate bias is calculated as in (7). Panel (b) shows the simulated tax bias as a function of third-party reported income, z, for tax evading wage earners and recipients of benefits ($\sim$194,000 obs.). Tax rate bias is calculated as in (6). In both panels, the dotted lines give the local average of the observations together with 95 percent confidence bands using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. The local mean smoothing in Panel (a) does not account for the stratification scheme. The simulated data in Panel (b) is not stratified. Income is defined as the sum of all income less deductions and is measured in '000 DKK. 1 USD $\approx$ 6 DKK (in 2006). In Panel (b), the budget is allocated such that approximately 4.1 percent of all wage earners and recipients of benefits are audited.

Figure 9, shows the corresponding figure of the data points and local averages of effective tax rates and residual income. Panel (a) and (b) share the same overall shape, namely, effective tax rates relatively unbiased (flat) in negative residual income but strongly biased and increasing in positive residual income. In both panels, the tax bias seems to decrease slightly at very high positive residual incomes. For the simulations, this reflects, similar to Figure 5(a) and 6(a), that high residual income is more common among taxpayers that also have large third-party reported incomes and that are audited relatively intensely. The structure of the data also seems consistent with this explanation. Again, local averages of effective tax bias in data and simulations in the distribution of residual income are highly correlated. The correlation, however, is somewhat smaller than for the progressive bias, reflecting the fact that regressive bias is generated partly by the allocation of audit probabilities within audit groups which, in the data, we observe imperfectly.

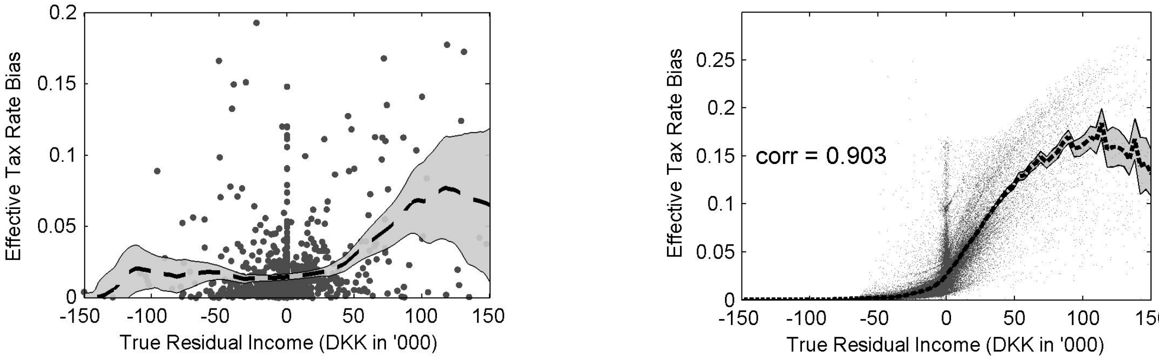

(a) Observed Tax Bias for Tax Evaders Across the Distribution of True Residual Income.

(b) Simulated Tax Bias for Tax Evaders Across the Distribution of True Residual Income.

Figure 9. Observed and Simulated Regressive Bias in True Residual Income.
Notes: The effective tax rate bias, $\tau - \tau^{\text{eff}}$, is the difference between the average statutory tax rate and the average effective tax rate as implied by the tax system, tax enforcement, and tax evasion behavior. Panel (a) shows the observed tax bias as a function of true, i.e., post-audit, income, u, for the subsample of tax evading wage earners and recipients of benefits (900 obs.). Tax rate bias is calculated as in (7). Panel (b) shows the simulated tax bias as a function of true, i.e., post-audit, income, u, for tax evading wage earners and recipients of benefits (~194,000 obs.). Tax rate bias is calculated as in (6). In both panels, the dotted lines give the local average of the observations together with 95 percent confidence bands using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. The local mean smoothing in Panel (a) does not account for the stratification scheme. The simulated data in Panel (b) is not stratified. Income is defined as the sum of all income less deductions and is measured in '000 DKK. 1 USD $\approx$ 6 DKK (in 2006). In Panel (b), the budget is allocated such that approximately 4.1 percent of all wage earners and recipients of benefits are audited.

4.3.3 Regressions

Another way to assess the correlation structure in the data is to run reduced-form regressions, as we have done in Table 4. First, Panel A shows estimates from running a median regression on the sample of tax evaders of tax evasion on true residual income, u, and third-party reported income, z, as well as a median regression of evasion on true total income, y. We allow slope coefficients to differ depending on whether u is positive or negative as we can see from Figure 7 that evasion behavior is markedly different for positive and negative true residual incomes.[27] Whereas evasion does not appear to be increasing in total income y, it is, in fact, strongly increasing (0.381 [0.050] for wage earners 0.029 [0.008] for the self-employed) in positive residual income. As we can see from

[27]In the model, this difference in evasion behavior is due to i) the incentive compatibility constraint dictates higher audit probabilities at lower reported residual incomes, which discourages evasion on negative residual income relative to positive residual income, and ii) the large mass point of honest tax reports at $x = 0$ combined with the incentive compatibility constraint necessitates low audit probabilities on the positive side, which encourages evasion on positive residual income relative to negative residual income.

Figure 7, this is because positive residual income is much easier to disguise – tax evaders for whom $u > 0$ simply evade the entire amount of their residual income. For the tax agency, these taxpayers are indistinguishable from the many honest taxpayers reporting around $u = 0$, so this type of evasion is costly to uncover. For wage earners this is an especially attractive strategy due to the large mass of honest taxpayers reporting $x = 0$. For the self-employed there is virtually no excess mass of honest taxpayers reported $x = 0$, but it is still the case that the conditional distributions of residual income given third-party reported income is centered around $u = 0$ which makes such a reporting strategy attractive. As we also noted above, Table 4 suggests that the observed average marginal propensity to evade taxes is smaller for the self-employed than for wage earners. In our model, this is explained by the much higher audit rate for self-employed compared to wage earners. Because the self-employed on average have higher incomes and are subject to less third-party reporting, a self-employed taxpayer would tend to evade more than a wage earner for the same audit risk. Despite the high audit rate for self-employed, they nevertheless evade substantially more than wage earners.

Next, Panel B shows marginal effects, multiplied by a factor 100 for readability, from a Tobit regression of audit flag intensity (our empirical counterpart to the audit probability) on third-party reported income, z, and reported residual income, x, allowing slopes to differ depending on whether x is positive or negative, and a Tobit regression of audit flag intensity on total reported income, $\tilde{y}$.[28] As the audit intensity is a function of not only reported residual income and third-party reported income, but also the distribution of true residual income at a given level of third-party reported income, we also include a function of u.[29] We find that the correlations exhibited by SKAT's audit flags are consistent with Scotchmer's conjecture that a population-wide description of an optimal enforcement regime should entail decreasing audit probabilities within audit groups but

[28] Left and right censoring of the audit flag intensity in the Tobit regressions is at 0 and 1, respectively.

[29] In fact, as we can see in Figure 6, in the aggregate for both observed and simulated data there is a slight positive correlation between the audit intensity and reported residual income for large/positive values of reported residual income. In the model there is a negative relationship between audit intensity and reported residual income within audit groups. The increasing relationship in the Figure 6(b) is caused by intense auditing of individuals with high levels of third-party reported income who also, on average, report relatively high residual incomes. The empirical distribution of audit intensity on third-party reported income in Figure 5(a) is consistent with this allocation of resources.

increasing probability of audit between groups, exhibited by the negative coefficient on $x \cdot D_{x \leq 0}$ and a positive coefficient on z.

The correlation of audit probability and income within and between audit groups translates into a significant positive correlation between audit probability and total reported income, $\tilde{y}$. Thus, despite the decreasing relationship within audit groups, third-party information reporting allows SKAT to audit taxpayers with high incomes more frequently.[30]

Lastly, in Panel C of Table 4 we run a median regression for the effective tax rate bias (in percentage points) either on true residual income, u, with slopes allowed to differ on the positive and negative domain of u, and third-party reported income, z, or on true total income, y. Since the data on tax evasion and audit flags seem consistent with the mechanisms driving the theoretical prediction of regressively biased effective average tax rates within audit groups, it is not surprising that we find a regressive bias within audit groups (i.e., a positive coefficient on $u \cdot D_{u \geq 0}$) and progressive bias between audit groups (i.e., a negative coefficient on z). These effects combine to make tax rates *progressive* in total income, y. As shown above, the progressive bias between audit groups derives from the fact the SKAT intensitvely audits taxpayers with high z.

Overall, Table 4 suggests a correlation structure of effective tax rates as depicted in the stilized Figure 1(b). The data supports the theoretical prediction that effective tax rates are regressive within audit groups. Between audit groups, there is a progressive bias such that average tax rates are actually progressively biased in total total income.

4.3.4 Non-Detectable Income

A potential problem for the robustness and validity of our results concerns non-detected tax evasion. As we discuss in Section 3, some unreported income is almost certainly missing from our measures of tax evasion, despite SKAT's dilligent effort in making

[30]The intention to audit high-income taxpayers with higher probability is not a specific feature of Danish tax enforcement. Internal Revenue Service (2012) shows how, in 2011, 1.0 percent of taxpayers with incomes less than \$200,000 were audited, 3.9 percent of taxpayers with income in the range of \$200,000-1,000,000 were audited, and 12.5 percent of taxpayers with incomes over \$1,000,000 were audited.

	Wage earners						Self-employed	
Dependent variable:	**A.** Evasion		**B.** Audit flag intensity		**C.** Tax bias (in pct. points)		**D.** Evasion	
$u \cdot D_{u\leq 0}$	-0.007	-	-0.040 **	-	-0.004	-	0.004	-
	(0.008)		(0.012)		(0.004)		(0.005)	
$u \cdot D_{u>0}$	0.381 ***	-	-0.000	-	0.030 ***	-	0.029 ***	-
	(0.050)		(0.003)		(0.003)		(0.008)	
z	0.000	-	0.002 **	-	-0.003 ***	-	0.006 ***	-
	(0.001)		(0.001)		(0.000)		(0.003)	
y	-	0.002	-	-	-	-0.002 ***	-	0.015 ***
		(0.001)				(0.000)		(0.004)
$x \cdot D_{x\leq 0}$	-	-	-0.034 **	-	-	-	-	-
			(0.011)					
$x \cdot D_{x>0}$	-	-	0.002	-	-	-	-	-
			(0.003)					
$\tilde{y}$	-	-	-	0.003 ***	-	-	-	-
				(0.001)				
Estimation method	Median regression		Tobit regression		Median regression		Median regression	
Constant term included	x	x	x	x	x	x	x	x
$D_{u>0}$ included	x		x		x		x	
$D_{x>0}$ included			x					
Sample	Evaders	Evaders	Full	Full	Evaders	Evaders	Evaders	Evaders
Observations	905	905	10,584	10,584	900	900	2,980	2,980
Obs. left-censored	-	-	8,555	8,555	-	-	-	-
Obs. right-censored	-	-	2	2	-	-	-	-
Pseudo R-squared	0.19	0.00	0.17	0.02	0.16	0.02	0.04	0.02

Table 4. Evasion Behavior, Tax Enforcement, and Tax Bias – Regressions on Data Sample.

Notes: *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.
All regressions are for wage earners and recipients of benefits except the rightmost two columns, which are regressions for self-employed. All samples are restricted to only contain taxpayers with strictly positive true total income.
Variable definitions: Monetary variables are in thousands ($u, z, y, \tilde{y}, x$, and evasion). Audit flag intensity is defined as no. of flags on the tax return divided by maximum no. of flags on any tax return (13) so that audit flag intensity is distributed on the unit interval. Tax bias is the percentage point difference between the true average tax rate and the effective average tax rate (reflecting both evasion behavior, audit probabilities, and penalty rates).
Evasion regressions (wage earners): Numbers in parentheses are stratified bootstrapped standard errors. The sample contains wage earners and benefit recipients with detected tax evasion and positive true income.
Audit flag regressions: Left and right censoring at 0 and 1, respectively. Numbers in parentheses are robust stratified standard errors. Estimates presented in the Tobit regressions are marginal effects. Marginal effects and standard errors in the Tobit regressions are multiplied by a factor of 100 for readability. The sample contains wage earners and benefit recipients with positive true income.
Tax bias regressions: Numbers in parentheses are stratified bootstrapped standard errors. The sample contains wage earners and benefit recipients with detected tax evasion and positive true income. There are five less evaders than in the evasion regressions. This is due to missing values in the compliance rating used to construct the tax bias measure.
Evasion regressions (self-employed): Numbers in parentheses are bootstrapped standard errors. The sample contains self-employed taxpayers with detected tax evasion and positive true income. As we do not have credible information on tax enforcement for the self-employed, there are no regressions for audit flag intensity and tax bias for the self-employed.

audits comprehensive. In particular, black market income is likely hard to detect. In Appendix B.1, we briefly present the best available evidence on the distribution of black market income in the Danish population based on survey data collected by the Rockwool Foundation Research Unit. This evidence suggests that black market income may be of a nonnegligible magnitude averaging approximately 3,143 DKK in the population. In comparison, the population weighted average underreported income is 3,619 DKK, cf. Table 2, Panel D.

However, assuming that black market income is completely non-detectable, the presence of such income will not affect neither the taxpayers' nor the tax agency's optimization criteria.[31] Consequently, the equilibrium of the model is unaffected, and the calibration exercise in this section remains valid because we fit the model to average tax evasion not including black market income.

Of course, even if black market income is completely non-detectable, it implies a measurement error in true residual income and translates into an underestimation of the effective average tax rate bias. Given that black market income is negatively correlated with reported income (cf. Appendix B.1) and third-party reported income is a very large part of reported income (approximately 95 percent in the population), we can deduce that black market income is also negatively correlated with third-party reported income. Therefore, accounting for black market income implies a level shift in the effective tax rate bias as a function of third-party reported income and, in addition, that this level shift is largest for taxpayers with little third-party reported income. As a result, including black market income implies a stronger progressive bias.

With respect to the regressive bias, the effect of including black market income depends on how black market income and detectable residual income are correlated. For the group of "wage earners", we know that low-income earners and recipients of public transfers more frequently provide black market labor. These individuals have little to no detectable residual income, and, as black market income constitutes positive residual income and at the same time increases the effective tax rate bias, including black market

[31] For taxpayers, this hinges on the assumption of risk neutrality. For example, if taxpayers are risk averse, non-detectable black market income may interact with tax evasion behavior.

income will tend to strengthen the positive relationship between tax rate bias and residual income (i.e., the regressive bias) for positive residual incomes depicted in Figure 9.

All in all, black market activities strenghten the distortions of the statutory tax system already generated by tax evasion and enforcement with respect to the formal economy.

4.3.5 Sensitivity Analysis

To show that our conclusions are robust to changes in parameters, we present in Figure 10 the simulated results of parameter changes for wage earners. We do this by changing the key parameters t_z, θ, and Q, and for each permutation letting $\mathbf{B}$ be calibrated to match simulated and observed average tax evasion among evaders. This we do for 27 permutations of the key parameters, i.e., all combinations of $-10\%, 0\%, +10\%$ changes to the set of parameters.[32] Focusing on the mechanism driving the regressive bias within audit groups, Panels (a) and (b) of Figure 10 show local averages of audit probabilities as a function of reported income and local averages of tax evasion as a function of true residual income, respectively. For audit probabilities, the changes are relatively minor, the main effects being a level shift in the maximal audit probability corresponding to changes in θ. For tax evasion, the local averages are all qualitatively similar, although the impact of parameter changes are larger among tax payers with larger residual incomes.

The structure of tax rate bias within and between audit groups in the simulations is also highly robust. The progressive bias between audit groups, shown in Panel (d), is virtually unchanged as it is generated mainly by the distribution of audit resources in the population, which is more or less unchanged by the parameter changes. The impact of parameter changes on the regressive bias within audit groups, shown in Panel (c), is more substantial as it compounds the effects of parameter changes shown in Panels (a) and (b). However, in all cases the qualitative relation between effective tax rate bias is very similar to the baseline simulation.

Varying the model parameters also affects the correlations of local averages in data and simulations, although not to a large extent. For the relationship between residual

[32] For t_z the changes are implemented as across-the-board increases/decreases in the marginal tax rate.

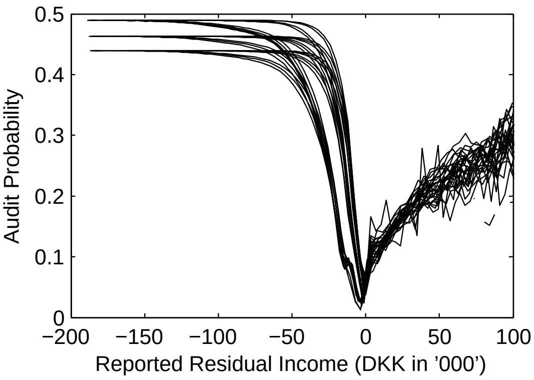

(a) Audit Probability, Reported Residual.

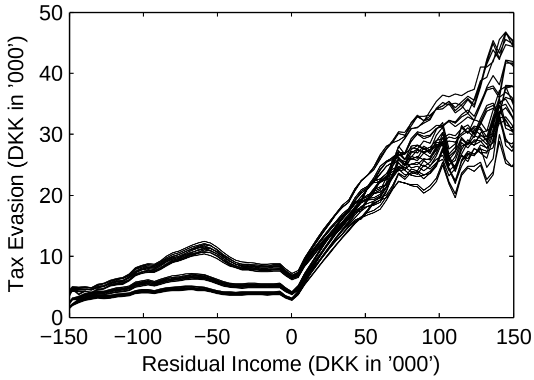

(b) Tax Evasion, True Residual Income.

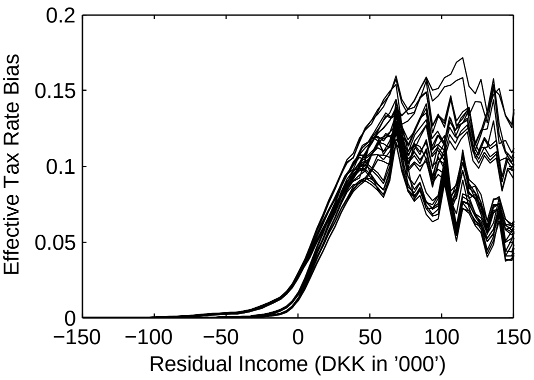

(c) Effective Tax Bias, True Residual Income.

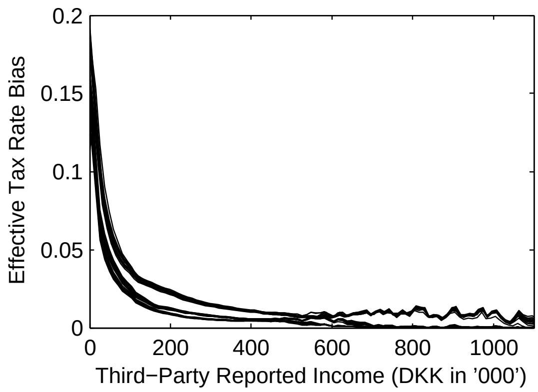

(d) Effective Tax Bias, Third-Party Reported Income.

Figure 10. Robustness Checks for Simulations of Wage Earners.
Notes: This figure checks the robustness of our simulation results graphically by plotting variations in estimated local average means on the basis of simulations with parameter permutations. We simulate the changes for wage earners in the four key relationships of the model, (a) audit probability as a function of reported residual income, (b) tax evasion as a function of true residual income, (c) regressive tax bias within audit groups as a function of true residual income, and (d) progressive tax bias between audit groups as a function of third-party reported income. The local averages depicted in the four panels are calculated in a similar manner to Figures 5–9, using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. We simulate the model for 10 percent parameters variations around the baseline estimates of t_z, θ, and Q, corresponding to 27 separate simulations. Thus, Panel (a) depicts variation in local means around the baseline simulation depicted in Figure 6(b) and similarly Panels (b)–(d) correspond to variations around the baseline simulated local means depicted in Figures 7(b), 9(b), and 8(b), respectively. All amounts in ’000 DKK (1 USD $\approx$ 6 DKK in 2006).

income and tax evasion, for example, the correlation coefficient lies between 0.875 and 0.952 compared to the baseline of 0.951. The most variable correlation the relationship between third-party reported income and the probability of audit which lies between 0.602 and 0.877 compared to the baseline of 0.803. However, this largely does not affect the correspondence of the progressive bias relationships between data and simulations – the correlation of local averages for third-party reported income and effective tax rate bias lies between 0.929 and 0.962. The regressive bias relationship varies more as it is affected by changes in both tax evasion and the audit probabilities within audit groups and lies between 0.666 and 0.907.

5 Concluding Remarks

This paper highlights the importance of information in tax enforcement. We find evidence in favor of the regressive bias prediction and Scotchmer's (1987) conjecture that it is crucial to distinguish regressive bias *within* an audit group from *aggregate* or *between*-group variation. Using detailed administrative data, we find evidence suggesting that, whereas effective tax rates are regressively biased within audit groups as theory suggests, this relationship is negated by a progressive bias between audit groups induced by the distribution of audit resources and third-party information. The outcome is that tax rates are progressively biased in total income. However, the model also suggests that an enforcement regime with much less third-party reported information would be substantially more regressive. In Denmark, this is avoided by a large information collection effort.

However, as emphasized by the literature, distortions may be substantial in settings in which third-party reporting is less comprehensive. The standard optimal auditing literature seems to suggest that regressive bias can be countered simply by adjusting marginal tax rates across the board. However, once we allow for population heterogeneity of behavior and income composition, this is no longer feasible. Our results suggest an obvious policy to ameliorate these distortions: increasing the share of income reported by third parties will reduce both the extent of evasion and the regressive bias in tax enforcement.

From a theoretical point of view, including third-party reported information and the likelihood of honest reporting conditional on the income composition is crucial in understanding tax evasion. A large literature on compliance versus non-compliance emphasizes behavioral/social explanations such as guilt and shame (e.g., Grasmick and Bursick, 1990), fairness (e.g., Spicer and Becker, 1980), and trust in government (e.g., Slemrod, 2003; Torgler, 2003). We focus on the implications, rather than the explanations, of honest reporting by some taxpayers. We analyze whether the observed moments and correlation structures of data are consistent with that generated by an optimizing tax agency, a group of honest taxpayers, and a group of expected utility maximizing tax evaders. We find that our model can replicate the extent of observed evasion as well as the subtle correlation structure of tax evasion, the probability of audit, and effective average tax rates with income. In addition, our results indicate that the Danish tax agency employs a distribution of resources across audit groups that is surprisingly similar in key respects to the optimal distribution generated by the model. All in all, there seems to be a role for both standard economic theory and behavioral/social extensions in explaining the behavior of tax evaders.

The correlation structure of effective tax rates seems robust: it is generated by our realistically complex model as well as in Scotchmer (1987). Furthermore, while variations in parameters change the level of average tax rate bias as well as the rate of progressivity between audit groups, in no variations is the correlation structure of effective tax rates qualitatively different from our baseline simulation. Thus, we are confident that similar empirical relationships would be found in data from any tax agency that employs, as does SKAT, a strong signal in predicting true incomes.

Based on data on the distribution of black market income in Denmark, we argue that our results are also robust to the lack of non-detectable income in our data. In fact, as discussed in Section 4.3.4, the data on the distribution of black market income suggests that our finding of regressive and progressive tax rate bias within and between audit groups, respectively, are lower bounds on the actual distortions of the statutory tax system. Under the assumption that black market income is non-detectable, the tax evasion and enforcement equilibrium is unaffected by the presence of this type of income, and the simulation results remain a valid description of optimal tax evasion and enforcement with respect to the formal economy. Moreover, the close fit between simulated and observed tax evasion and enforcement indicates that the assumption of complete non-detectability of black market income is an appropriate simplification.

A natural objection to the model we employ is the lack of general interactions with labor market choices. We accomodate to some extent the self-selection of taxpayers into employment categories by allowing the fraction of honest taxpayers to differ between wage earners and self-employed. In addition, disregarding dynamic aspects is not likely to be important due to the limited retrospectivity of SKAT's actual audit scheme and the restrictive statute of limitations on retroactive penalties for tax evasion. However, we do not account for other effects, e.g., how tax enforcement affects labor supply on the intensive margin.

Despite these limitations, our paper advances the literature in the direction of developing a full-fledged structural model of tax evasion, which can be estimated directly with maximum likelihood or GMM methods. Moreover, our paper is an important next step towards an understanding of the comparative statics of tax evasion and enforcement. Such an avenue of research may be seen as a necessary complement to the literature on their behavioral and social determinants, which have been extensively explored in the literature.

Appendix

A.1 Numerical Implementation

The second order differential equation is obtained by combining (1), (3), (4), and the expression for $\frac{\partial u}{\partial x}$ to get

$$p''(x) = \left(\frac{f_u^h(x)\, \frac{\lambda c}{\theta t + t}}{f_u^d(u(x)) \left[\frac{p(x) - \frac{1}{\theta+1}}{p'(x)} - \frac{\lambda c}{\theta t + t} \right]} - 2 \right) \cdot p'(x)^2 \left(\frac{1}{1+\theta} - p(x) \right)^{-1}, \tag{A.1}$$

suppressing z for convenience. Thus, sufficient conditions for equilibrium, given $B(z)$, are the two equations (5) and (A.1).

We approximate the equilibrium solution by discretizing z into an evenly spaced grid point vector of dimension 200.[33] Equilibrium functions for other values of z are approximated by interpolation. For each gridpoint, we solve the 2nd order ordinary differential equation (ODE) in (A.1) for many values of $\overline{x}$, where $\overline{x} \equiv x(\overline{u})$. The ODE algorithm is then initialized using $p(\overline{x}) = 0$ and $p'(\overline{x}) = \left(\frac{1}{1+\theta}\right) / (\overline{u} - \overline{x})$, cf. (1). For each value of $\overline{x}$ and z, we need a corresponding value of $\lambda(z)$, the shadow value of increasing the budget size. However, $\lambda(z)$ and $\overline{x}$ are not separately identified. Therefore, we must take a heuristic approach, solving for each $\overline{x}$ the ODE for many values of λ until one is found that satisfies the equilibrium conditions everywhere, in particular at $x = \overline{u}$. In practice, we do not merely guess repeatedly at $\lambda(z)$, but employ a search algorithm to find the $\lambda(z)$ that satisfies (5); this provides a candidate $\lambda(z)$ corresponding to a particular $\overline{x}$ that satisfy the FOC everywhere with a small error tolerance. Figures 1(a) and 1(b) illustrate an example of the set of solutions resulting from the algorithm.

When this algorithm has executed for all grid points of z, we can determine the optimal budget allocation using the fact that in an interior equilibrium $\lambda(z)$ must be equalized across different levels of z.

[33] The model for the self-employed is substantially more computationally intensive so there we use only 100 grid points. Of course, this implies that interpolations will be less precise, but this does not appear to be important. Likewise, solutions using only 50 grid points are graphically indistinguishable in terms of Figure 7.

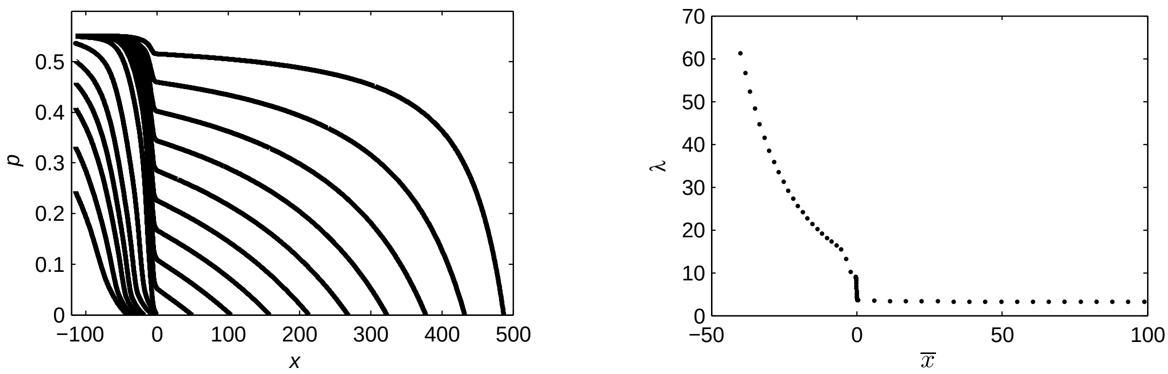

(a) Examples of Optimal Audit Functions: $p(x|z)$.

(b) Shadow Values with respect to $B(z)$.

Figure A.1. Solutions Examples
Notes: $\bar{x}$ (xbar) is defined as the lowest value of x that solves $p(x|\cdot) = 0$, i.e., the highest report of dishonest taxpayers. Reported residual, x, is measured in '000 DKK (1 USD $\approx$ 6 DKK in 2006).

Equation (A.1) can be solved by standard numerical methods. We employ a Runge-Kutta-type algorithm developed in Shampine (2009), which outperforms most standard ODE algorithms in terms of precision and robustness. However, two main numerical issues must be resolved.

First, due to point mass in $f^h_{u|z}$ at $u = 0$, $\mathrm{E}(u|x, z)$ is discontinuous at $x = 0$, which induces what is known as a "singularity" in the differential equation. We take a standard approach to this problem and approximate solutions for which $\overline{x} > 0$ by substituting the logical function $\mathbf{1}_{(x=0)}$ with a smooth, differentiable approximation. The resulting function displays a relatively smooth transition from 0 to 1 in a small band around $x = 0$. An alternative approach is to split the ODE algorithm in two, corresponding to the domains $[\underline{u}, 0)$ and $[0, \overline{x}]$, and identifying the discontinuous jump in $p'(x)$ from the equations characterizing the equilibrium and the measure of point mass at $x = 0$. However, as the size of this discontinuity cannot be identified analytically, this introduces an element of imprecision in the algorithm which, in our experience, may negatively affect the robustness of the algorithm.

Second, the ODE algorithm may fail to converge if we allow the conditional density function to take values extremely close to 0 since the ratio $\frac{f^h_{u|z}(x)}{f^d_{u|z}(u(x))}$ may diverge toward infinity. Estimating the density f_{uz} as a bivariate kernel density is numerically inconvenient as it tends to result in "troughs" of zero density in the interior of the domain of

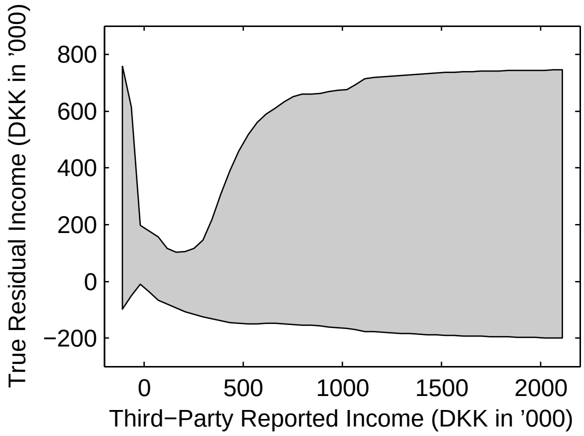

Figure A.2. The Support of u Across Audit Groups.
Notes: The estimated conditional densities of $u|z$ for wage earners and benefit recipients are truncated at the 0.25 and 99.75 percent fractiles of the unrestricted conditional distributions. Residual income, u, and third-party reported income, z, are measured in '000 DKK (1 USD $\approx$ 6 DKK in 2006).

some conditional distributions. Instead, as mentioned in the main body of the paper, we estimate f_{uz} as a lognormal mixture distribution.[34] Specifically, the mixture distribution consists of six component distributions. Increasing the number of component distributions allows a more flexible fit of the distribution but alters our results only very slightly. Lastly, we truncate the domain of the potential tax evaders' conditional true income distributions where the densities are very close to zero to keep the fraction in equation (A.1) from diverging to infinity. Specifically, we truncate the unrestricted conditional densities at the 0.25 percent and 99.75 percent fractiles. The resulting supports of the conditional distributions vary in z as illustrated in Figure A.2.

B.1 Black Market Activities

A potentially important avenue for tax evasion is black market income. This type of income is much harder to discover by tax auditors and, thus, less likely to be included in our data despite the intensive auditing of tax returns for the experiment. To quantify the extent of black market activities in the population at large, we utilize survey data

[34]Of course, the lognormal distribution is not defined on domains that include negative values. In practice, in estimating the mixture distribution, we create a simple additive mapping of the observations to a set of "virtual residual incomes" that are entirely positive, estimate the lognormal mixture distribution using six component distributions, and use the mapping to obtain the actual bivariate income distribution. The resulting distribution is indeed very close to that obtained by using a bivariate kernel density algorithm.

collected by the Danish Rockwool Foundation Research Unit. They have since 1985 collected survey data for random samples of the population in an attempt to quantify the incidence of and return to black market activity. Although the surveys in principle collect identifiable information, such as social security numbers, the data set is anonymized. As such, no one has ever been prosecuted for having black market income due to participation in these surveys. Unfortunately, there were no surveys carried out in 2006 so we use two surveys from 2005 and two from 2007 instead (all amounts in 2006-prices). The surveys contain many variables, but we focus on measures of the incidence of and return to the supply of black market services (i.e., remuneration for black market labor). In Figure B.1, we show the incidence of black market work and the return to this activity across the distribution of reported income. Panel (a) shows the share of taxpayers having performed black market work for 20th fractiles of the reported income distribution in the sample. The figure indicates that black market work is more common among low-income taxpayers, whereas middle and top earners figure less prominently.[35] Panel (b) shows the average black market income across the distribution of reported income for the entire sample, unconditional on whether or not taxpayers have participated in black market work, using local mean smoothing. Again, mainly low-income taxpayers have black market labor income. This is consistent with a comprehensive study by the Danish Economic Council in 2011 (DØRS, 2011) using the same data as here but for all available years, which concluded that black market earnings were negatively correlated with total reported income. In addition, DØRS (2011) finds that the self-employed and low-income wage earners most frequently supply labor on the black market, and black market wages are substantially higher for the self-employed, indicating that this group of taxpayers on average earns larger black market incomes than wage earners.

[35] In the survey samples from 2005 and 2007 that we use, there are very few top earners. But in DØRS (2011), which uses a larger sample spanning more waves of the survey, there is a clear picture that top earners supply black market work less frequently than the middle income earners.

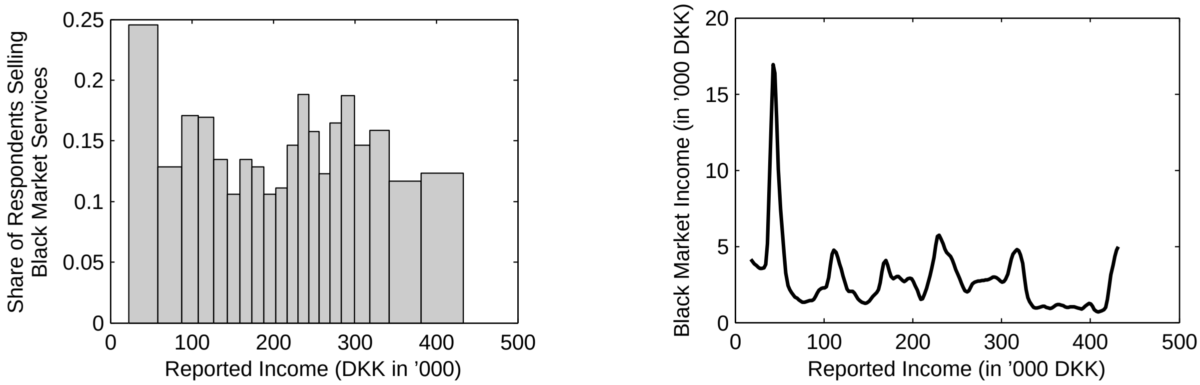

(a) Share of Taxpayers Selling Black Market Services by 20th Fractiles of the Reported Income Distrition.

(b) Mean Black Market Income Across the Reported Income Distribution.

Figure B.1. Size and Distribution of the Black Market Economy in Denmark.
Notes: The data on black market activity stems from a survey collection effort undertaken by the Rockwool Foundation Research Unit, and income data stems from linked administrative data. The Rockwool Foundation Research Unit's surveys on black market income has been collected since 1985. For each wave in the collection effort, surveys are dispatched to individuals, both wage earners and the self-employed, with the understanding that their answers are kept anonymous. As such, no one has ever been prosecuted for acknowledging black market income unreported on their tax return in the surveys. Unfortunately, the surveys were not collected in 2006. Instead, we have obtained data for four surveys collected in 2005 and 2007 (in 2006 prices), for a total sample of surveyed individuals of 3,806. Only 10 individuals did not respond so the sample of responsive individuals is 3,796. Of these individuals, 560 responded that they had sold services on the black market during the last 12 months with an average income of 19,439 DKK. In the total sample of responsive individuals, this corresponds to an average black market income of 3,143 DKK. Panel (a) shows the share of taxpayers in the sample selling black market services by 20th fractiles of the distribution of total reported income. Panel (b) shows the average return to black market activity in the sample using local mean smoothing with the Epanechnikov kernel function and a rule-of-thumb bandwidth. All amounts are in '000 DKK (1 USD $\approx$ 6 DKK in 2006).
Source: Rockwool Foundation Research Unit and own calculations.

References

Allingham, M. G., Sandmo, A., 1972. Income tax evasion: A theoretical analysis. Journal of Public Economics 1, 323–338.

Alm, J., McClelland, G. H., Schulze, W. D., 1992. Why do people pay taxes? Journal of Public Economics 48 (1), 21–38.

Andreoni, J., Erard, B., Feinstein, J. S., 1998. Tax compliance. Journal of Economic Literature 36 (2), 818–60.

Clotfelter, C., 1983. Tax evasion and tax rates: An analysis of individual returns. The Review of Economics and Statistics 65 (3), 363–373.

Cremer, H., Gahvari, F., 1994. Tax Evasion, Concealment and the Optimal Linear Income Tax. Scandinavian Journal of Economics 96 (gr), 219–239.

Cremer, H., Marchand, M., Pestieau, P., 1990. Evading, auditing and taxing: The equity-compliance tradeoff. Journal of Public Economics 43 (1), 67–92.

DØRS, 2011. Dansk økonomi, Forår 2011. www.dors.dk.

Erard, B., Feinstein, J. S., 1994. Honesty and evasion in the tax compliance game. RAND Journal of Economics 25 (1), 1–19.

Feinstein, J., 1991. An econometric analysis of income tax evasion and its detection. The RAND Journal of Economics 22 (1), 14–35.

Feld, L. P., Frey, B. S., 2002. Trust Breeds Trust: How Taxpayers Are Treated. Economics of Governance 3 (2), 87–99.

Grasmick, H. G., Bursick, Jr., R. J., 1990. Conscience, significant others, and rational choice: Extending the deterrence model. Law and Society Review 24 (3), 837–861.

Internal Revenue Service, 2012. Fiscal year 2011 enforcement and service results. http://www.irs.gov/newsroom/article/0,,id=251923,00.html.

Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., Saez, E., 2011. Unwilling or unable to cheat? Evidence from a randomized tax audit experiment in Denmark. Econometrica 79 (3), 651–692.

Nielsen, S. B., Sørensen, P. B., 1997. On the optimality of the nordic system of dual income taxation. Journal of Public Economics 63, 311–329.

Phillips, M. D., 2010. Taxpayer Response to Targeted Audits. mimeo, University of Chicago.

Pomeranz, D., 2010. No taxation without information: Deterrence and self-enforcement in the value added tax.

Reinganum, J. F., Wilde, L. L., 1986. Equilibrium verification and reporting policies in a model of tax compliance. International Economic Review 27 (3), 739–60.

Sanchez, I., Sobel, J., 1993. Hierarchical design and enforcement of income tax policies. Journal of Public Economics 50 (3), 345–69.

Scotchmer, S., 1987. Audit classes and tax enforcement policy. American Economic Review 77 (2), 229–33.

Scotchmer, S., 1992. The regressive bias in tax enforcement. Public Finance 47, 367–371.

Shampine, L. F., 2009. Vectorized solution of ODEs in Matlab with control of residual and error. `http://faculty.smu.edu/shampine/`.

Slemrod, J., 2003. Public Finance and Public Policy in the New Century. Cambridge, MA: MIT Press, Ch. Trust in Public Finance, pp. 49–88.

Slemrod, J., 2007. Cheating ourselves: The economics of tax evasion. Journal of Economic Perspectives 21 (1), 25–48.

Spicer, M. W., Becker, L. A., 1980. Fiscal inequity and tax evasion: An experimental approach. National Tax Journal 33 (2), 171–175.

Torgler, B., 2003. Tax morale, rule-governed behavior, and trust. Constitutional Political Economy 14 (2), 119–140.

Chapter 2

A Structurally Estimated Model of Tax Evasion and Enforcement*

Jori Veng Pinje and Simon Halphen Boserup

Department of Economics, University of Copenhagen

November 2012

Abstract

We set up and structurally estimate a model of tax evasion and optimal auditing. The model builds on both rational tax evasion, the well-established fact that some taxpayers are inherently honest, and takes into account that modern tax collection relies on information reporting in addition to traditional auditing. Leveraging a uniquely detailed data set of random audits in Denmark that allows us to control for the use of third-party reporting of income in the audit process, we assess the relative importance of instruments to deter tax evasion. We find that the policy instruments that work along the intensive margin of tax evasion (audit and penalty rates) are less effective in combating tax evasion than instruments working along the extensive margin of tax evasion (third-party information reporting and the share of honest taxpayers in the population).

*We are grateful to the Danish tax administration, SKAT, for providing data access. We thank Claus Thustrup Kreiner and David Dreyer Lassen for helpful comments and discussions. The authors can be contacted at jori.pinje@econ.ku.dk and simon.h.boserup@econ.ku.dk.

1 Introduction

In recent years of fiscal peril across the world, securing tax revenue has become ever more important. Likewise, the cost borne by tax authorities collecting taxes and tax filers spending hours are immense.[1] Despite this, the economic literature still faces severe empirical challenges when it comes to understanding the nature and extent of tax compliance and evasion. Even basic empirical relationships such as, say, the relationship between income and the propensity to evade have remained uncertain due to conflicting results; see, e.g., Clotfelter (1983) and Feinstein (1991).

As main predictions of the early expected-utility model of tax evasion of Allingham and Sandmo (1972) and variants thereof could not be substantiated empirically, the empirical literature largely turned towards investigations into behavioral deviations from the expected utility framework. Initially, the argument against models built on standard expected utility maximization relies on the apparent incompatibility of the relatively low levels of evasion observed with the low share of taxpayers actually audited (say, 1–2 percent), and the even lower share of taxpayers penalized for tax evasion, given reasonable degrees of risk aversion; see, e.g., Alm, McClelland, and Schulze (1992b) and Feld and Frey (2002).

However, as argued in Slemrod (2007), this ignores the fact that the average audit rate vastly underestimates the risk of an audit on the average dollar actually evaded. Tax authorities in developed countries make extensive use of information reporting from third parties. In the US, approximately 3/4 of all income is reported to the IRS by third parties (Andreoni, Erard, and Feinstein, 1998), and in Denmark, which this paper focuses on, that number is closer to 95 percent (Kleven, Knudsen, Kreiner, Pedersen, and Saez, 2011). This information makes certain types of income ill-suited for tax evasion purposes. For example, if employees' salaries are reported to the tax authorities by their employers, the risk of audit on unreported salary income is much higher and probably closer to 100 percent. However, because which line items are subject to third-party reporting

[1] Taxpayer Advocate Service, IRS (2010) estimates that overall tax compliance in the U.S. consumes 6.1 billion hours a year spent by taxpayers and businesses alone, equivalent to more than 3 million full-time workers.

and which are not is common knowledge to both tax agency and taxpayers, little to no underreporting occurs for these income components. Consequently, the high probabilities of audit as a consequence of evasion on line items subject to third-party reporting remains unobserved.

Following Slemrod (2007), a series of recent papers (Phillips, 2010; Pomeranz, 2010; Kleven et al., 2011; Boserup and Pinje, 2012) have re-emphasized the rational choice paradigm: in different ways, these papers stress the importance the informational structure of the tax evasion game. In particular, these papers argue that the case for the expected utility paradigm becomes much stronger once we account for tax agencies' abilities to leverage *ex ante* income information to target audits to likely evaders.

Although the initial critique of the Allingham-Sandmo framework has proved unwarranted, the behavioral literature has nonetheless shown that there is room for behavioral extensions to the classical framework. There is a large, mainly experimental, empirical literature on honesty, social norms, tax morale, belief formation, etc., and how these concepts are related to tax evasion (e.g., see Andreoni et al., 1998; Slemrod, 2007, for surveys.).

In this paper, we structurally estimate a game-theoretic model of tax evasion and optimal auditing containing elements from both the expected utility approach and the behavioral literature. The two most prominent features of the model are the presence of third-party information reports and a share of honest taxpayers in the population, both of which affect the degree to which the tax agency can infer tax evasion on the basis of tax returns. With this model we are able to account for the average level of tax evasion of wage earners in Denmark, and we use the estimated model to evaluate the relative importance of policy parameters working along the intensive and extensive margins of tax evasion. On the intensive margin, policy instruments affect the trade-off of expected costs and benefits from tax evasion on the margin. The audit rates and penalty rate of the enforcement system are such instruments. On the extensive margin, policy instruments affect the expected costs and benefits from engaging or not engaging in tax evasion. Policies that affect honesty or the degree of third-party reporting are such instruments as they entail truthful reporting on discrete amounts of income.

The contributions of this paper are three-fold. 1) To our knowledge, we are first to structurally estimate a model of tax evasion and optimal auditing. We do this using a very detailed data set from Kleven et al. (2011) containing individual level data on complete tax returns prior to and following intensive audits as well as the amounts of third-party reported and self-reported income by line item. 2) Further, we show, by incorporating the use of third-party reporting and a share of honest taxpayers into the model, that it is possible to generate the level of observed evasion for reasonable parameters using a model based on expected utility maximization. 3) Simulating the effects of an actual 2010-2012 Danish policy change, in which third-party reporting was introduced on capital gains from stocks, and revenue-equivalent counterfactual parameter changes, we find suggestive evidence that policies that work along the intensive margin of tax evasion are less effective than those working along the extensive margin of tax evasion.

We now proceed to the main body of the paper. Section 2 briefly outlines the model. Section 3 describes the empirical strategy for obtaining parameter estimates and Section 4 provides policy experiments. Section 5 concludes.

2 Model

We employ a model developed in Erard and Feinstein (1994) and adapted for empirical use in Boserup and Pinje (2012). Whereas several current theories are capable of analyzing behavior *within* audit groups, i.e., conditional on pre-defined groups based on *ex ante* observable information, this model allows us to simulate *aggregate* reporting behavior as well as the tax agency's audit response. To this end Boserup and Pinje (2012) generalize the model in Erard and Feinstein (1994) to incorporate a population that is heterogeneous in third-party income reports and where different levels of third-party reported income plays the role of audit groups.

The model's main feature is introduction of noise in taxpayer reports by incorporating the stylized fact that some taxpayers report their incomes honestly, even when they have ample opportunity to evade taxes. This increases the realism and the empirical applicability of the model because it eliminates several potential equilibria and leaves a unique revenue maximizing equilibrium prediction.[2] The resulting model provides a relatively simple framework for analyzing the behavior of rational tax evaders and the tax authority given the inherent informational asymmetries that stem from reporting type (i.e., honest/evading) and individual earned income unknown to the tax agency.

In the model, the tax agency selects the audit regime subject to a budget constraint without being able to commit to an audit strategy. The audit schedule for a particular audit group (i.e., conditional on a particular third-party reported income level) is a function of taxpayers' reported *residual* incomes, i.e., their income in excess of third-party reported income. The tax agency allocates its resources across different strata of the population so as to equalize the shadow values of extending resources to auditing taxpayers with different amounts of third-party reported income. Whereas the distribution of true incomes, conditional on information reports, is known, actual true incomes of individual taxpayers are private information. Taxpayers choose income reports subject to their expectations about the audit regime.

2.1 Individual Reporting Behavior

Individual taxpayers have true taxable incomes y and report taxable incomes, $\tilde{y}$. Part of true income, z, is reported by third parties and is known to all agents. Therefore, $y = z + u$, where u is residual income, which can be positive or negative as it includes both, e.g., wages and deductions not reported by third parties. u is *ex ante* unknown and can only be ascertained by the tax agency by conducting a costly audit, which is assumed to completely reveal true residual income. The *reported* residual, x, is $x = \tilde{y} - z$.

Taxpayers are split into three groups, those that can and will evade taxes (evaders), those that could but do not evade taxes, and those that cannot evade taxes (mass point taxpayers). In principle, of course, it is possible to evade taxes despite having zero residual

[2] In contrast, most other models contend with unrealistic predictions, multiple equilibria, or both. See Reinganum and Wilde (1985) and Border and Sobel (1987) which contend with multiple equilibria and, for example, develop equilibria in which the tax agency, in equilibrium, can infer true incomes of all individuals based on reports and rarely or never audits any actual tax evaders.

income, for example by claiming unwarranted deductions. However, of the observed taxpayers having zero residual income, very few taxpayers in our sample take advantage of this possibility. The reason for this is likely that evading taxes by exaggerating deductions requires some deductions to begin with. Below, we will denote potential tax evaders "evaders" and compliant taxpayers, either because they cannot or will not evade taxes, for "honest taxpayers".

For simplicity, taxes are assumed to be linear.[3] Whereas honest taxpayers always report $x = u$, we assume that dishonest taxpayers are risk neutral and maximize expected utility given by expected income net of taxes and penalties

$$(1-t)\,z + p\,(x|z)\,[(1-t)\,u - \theta t\,(u-x)] + (1 - p\,(x|z))\,[u - tx]\,,$$

where t is the tax rate, θ is the penalty rate on tax evasion, and $p(x|z)$ is the audit probability for report x given the level of third-party reporting z.[4] The correct amount of taxes are paid with certainty on income reported by third parties, whereas taxes (and penalties) paid on residual income depends on both the evasion behavior of the taxpayer and whether or not the taxpayer is audited.

In optimum, the taxpayer's choice must satisfy the first order condition

$$u = x + \frac{p\,(x|z) - \frac{1}{1+\theta}}{p'\,(x|z)}. \tag{1}$$

It is clear from Equation (1) that for $p\,(\cdot) = \frac{1}{1+\theta}$, $x = u$ and evasion is discouraged completely. However, $p \geq \frac{1}{1+\theta}$ is not compatible with equilibrium when the tax agency cannot commit to the audit regime: if evasion were completely discouraged, the tax agency would lower p for some x as a cost saving measure. Thus, in equilibrium $p\,(\cdot) \in \left[0, \frac{1}{1+\theta}\right)$.

[3] However, as described in Boserup and Pinje (2012), in the simulations we allow the constant marginal tax rate to vary by z to partly allow for progressiveness of the actual income tax scheme. For more information on the Danish tax system in 2006, see Boserup and Pinje (2012) and Kleven et al. (2011).

[4] Risk neutrality is perhaps an unrealistic assumption in this context. However, as noted in Boserup and Pinje (2012), due to the large extent of third-party reporting in the Danish tax system, simulations of risk neutral vs. risk averse taxpayers yield very similar results. We opt for using risk neutrality here because the numerical algorithm requires many evaluations and numerical algorithm for the model with risk averse taxpayers is much more computationally demanding.

Given that $p'(x|z)$ is negative and $p(x|z) < \frac{1}{1+\theta}$, increasing the audit probability will, *ceteris paribus*, lower tax evasion as the risk of getting caught is higher. Lowering $p'(x|z)$ (increasing its absolute value) also reduces tax evasion by increasing the risk of audit from taxes evaded on the margin.[5]

2.2 Optimal Audit Response

The tax agency chooses a continuum of audit schedules, $p(x|z)$, and a budget allocation, $B(z)$, for all z. In this way, the informational aspect of using third-party reported incomes to predict true income is incorporated into the population-wide equilibrium.[6]

Denote the densities of true income conditional on third-party reports $f^h_{u|z}$ and $f^d_{u|z}$ for honest and dishonest taxpayers, respectively. In addition, define the total density function as $f_{u|z} = f^h_{u|z} + f^d_{u|z}$ and $F_{u|z}$ the conditional distribution function associated with $f_{u|z}$.

The audit schedule is chosen to maximize expected revenue (taxes plus fines)

$$\int \left(\int_{\underline{x}}^{\bar{u}} \left[p(x|z) \left(t\mathrm{E}(y|x,z) + \theta t \left(\mathrm{E}(y|x,z) - \tilde{y} \right) \right) + (1 - p(x|z))\, t\tilde{y} \right] \mathrm{d}F_{x|z} \right) \mathrm{d}F_z \tag{2}$$

subject to the budget constraint

$$c \int \left(\int_{\underline{x}}^{\bar{u}} p(x|z)\, \mathrm{d}F_{x|z} \right) \mathrm{d}F_z \leq \int B(z)\, \mathrm{d}F_z \equiv \mathbf{B},$$

where $F_{x|z}$ is the induced conditional distribution function for reported residual income, x, given third-party reported income, z; F_z is the marginal distribution function for z; and $B(z)$ is the proportion or density of the overall audit budget, $\mathbf{B}$, allocated to income reports with third-party reported income, z. We normalize the cost per audit, c, to one.

[5] Taxpayers' income returns must also satisfy the second order condition, $p''(x|z)(x-u) + 2p'(x|z) \leq 0$.

[6] In principle, the tax agency could also condition audit schedules on other population variables such as gender, age, occupation, etc. However, as Kleven et al. (2011) show, these variables are less powerful as predictors. Conditioning on whether the taxpayer was audited in previous years would complicate matters as it would introduce a dynamic aspect to reporting decisions. However, as observations on past audits are not employed in the Danish tax authorities' actual audit scheme, this limitation is unlikely to affect the fit of our model. In addition, the statute of limitations for retrospective audits is limited to 14 months.

Consequently, $\mathbf{B}$ may be interpreted as the share of taxpayers audited. For each (x, z), the tax agency must choose p to solve

$$\max_{p} \{p [t\mathrm{E}(y|x,z) + \theta t (\mathrm{E}(y|x,z) - \tilde{y})] + (1-p)t\tilde{y}\} \,\mathrm{d}F_{x|z}\mathrm{d}F_z \\ -\lambda(z)\, c \left[p \;\mathrm{d}F_{x|z} - B(z)\right] \mathrm{d}F_z,$$

where $\lambda(z)$ is the Lagrangian multiplier on the budget constraint. From this condition, a unique mixed audit strategy can be found by expressing the expectation $E(y|x,z)$ in terms of conditional income densities and relationships between reporting behavior and true income. Utilizing the optimality requirement that the shadow values of increasing the budget size in an audit group must be equated across different values of z, we can express the equilibrium conditions as a second order differential equation which depends on the unknown constant, λ. Lastly, λ is identified by a boundary condition on the reporting behavior for taxpayers claiming residual incomes of $x = \underline{u}$, i.e., the lower bound of the residual income distribution. In equilibrium, some proportion of tax evaders pool their reports at $\underline{u}$, and the boundary condition requires that the tax agency's first order condition is satisfied for such reports.

The unique revenue maximizing equilibrium of the model is described by the collection of functions, $u(x|z)$ and $p(x|z)$, and the budget distribution, $B(z)$.

Figure 1(a) shows an example of what the audit schedule, $p(x)$, looks like: it starts at the lower bound of the residual income distribution, $\underline{u}$, is downward sloping, and terminates in $p(\bar{x}) = 0$. This form balances the need to audit in order to raise revenue with the cost of doing so. The negative slope reflects the need to discourage high-income taxpayers from reporting too low incomes.

Figure 1(b) shows the amount of evasion as a function of true income. The linear increase in the first part of the graph reflects pooling of dishonest taxpayers: for a given audit schedule, there will be some level of residual income, u^{pool} in $[\underline{u}, \overline{u}]$, for which the most profitable report is $\underline{u}$. Consequently, all taxpayers with residual incomes below u^{pool} also report $x = \underline{u}$. Therefore, there will be a point mass in the induced distribution

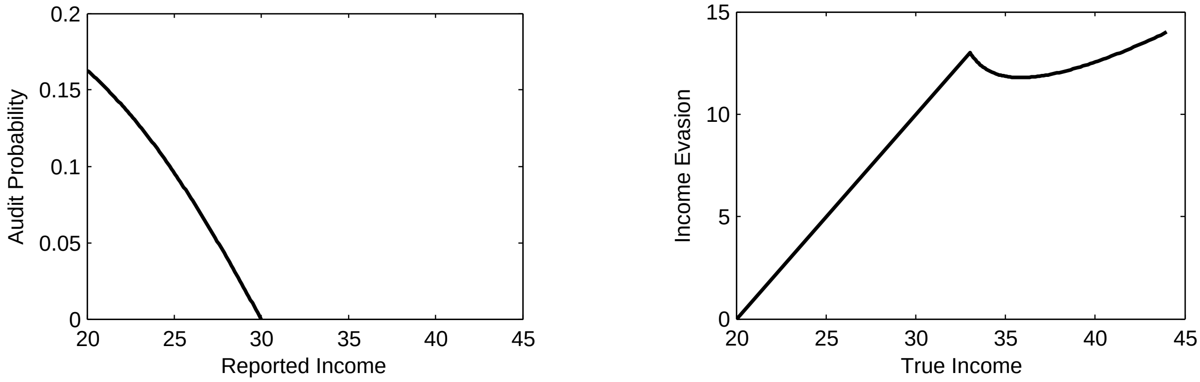

(a) The Optimal Audit Schedule, $p(x)$.

(b) Evaded Income, $u - x$, by True Income, y, for Dishonest Taxpayers.

Figure 1. Equilibrium Responses and Tax Bias.
Notes: Both panels display an example of equilibrium functions from the Erard and Feinstein (1994) model without third-party reporting. Equivalently, this could be an example of the solution for a particular z in our model including third-party reporting. This example is produced assuming $B = 10$ percent, $\log(u) \sim \mathcal{N}(3.42, 0.3^2)$ truncated on $[20, 44]$, $Q = 0.4$, and $t = 0.5$.

of reports, $f_x(x)$. After this pooling point, evasion falls rapidly in income until evasion again becomes increasing in income as the probability of detection becomes sufficiently low.

3 Empirical Strategy and Results

3.1 The Danish System of Filing Taxes

Taxpayers' incomes earned during some year t are reported to the tax authorities (SKAT) by entities such as employers, banks, pension funds by the end of January year $t+1$. In general, all income received as salary, private/public pensions, honorarium, unemployment benefits, etc., is subject to third-party reporting as is the case with, e.g., mortgage interest payments and some capital income; in 2006/07 dividends were but capital gains were not reported by third parties. Self-employment income is rarely covered by information reporting. Only in cases where, e.g., remuneration is paid by a public institution will such income be subject to third-party reporting. Third parties have no discretion as to whether or not to supply SKAT with this information. The requirement is entirely related to income type.

By mid-March, SKAT sends out "pre-populated" tax returns to taxpayers using whatever third-party reports they have received as well as other information (e.g., private and work addresses are used for calculating commuting allowances). Taxpayers then have until May 1 to correct the tax return. If the taxpayer does not change his or her tax return, the pre-populated return counts as final. After May 1, SKAT's computer system processes the tax returns and attaches (binary) audit flags to returns the system finds likely to contain errors. Afterward, tax examiners assess the flagged returns and decides, based on information available, local knowledge, and audit resources available, whether or not to initiate an audit. Each instance of an audit flag has to be actively closed by a tax examiner.

In short, the flag system relies on the third-party reports and auditing "best practices" that could be converted to algorithmic form, e.g., specific tax return compositions indicative of misreporting, cut-off rules based on expected incomes conditional on third-party reported income, etc. The flag system consists of a large number of flags, each of which is intended to signal the likelihood of tax evasion on particular line-items or combinations of line-items.

3.2 Data

Data stems from a large-scale compliance project carried out by the Danish tax authorities in 2006-08 and previously studied in Kleven et al. (2011) and Boserup and Pinje (2012). The compliance project comprised of a series of tax audit rounds and different treatments. We use data from the first round baseline audits.

Much like the U.S. Taxpayer Compliance Measurement Program, the baseline sample constitutes a stratified[7] random sample of taxpayers that were subject to extensive, homogeneous audits. We use the sample of 10,740 non-treated individuals that were either employees or recipients of public benefits (i.e., not self-employed), from the fiscal year 2006 as information reporting for this group is abundant. We denote these "wage earners". The full population of taxpayers in this group was approximately 4.2 million

[7] Wage earners are stratified on tax return complexity. See notes to Table 1 for further details.

in 2006.

After the deadline for changing one's tax return (May 1, 2007), SKAT carried out unannounced audits on the selected group. These audits were comprehensive covering all items on the tax return (in contrast to regular audits motivated by audit flags), demanding documentation for all items on which SKAT did not itself possess information, and cross-checking with other data sources whenever possible. SKAT made a significant effort to homogenize the audits by organizing training workshops, distributing detailed audit manuals, etc. The audits of all the wage earners and a group of self-employed taxpayers in the experiment took up 21 percent of SKAT's annual resources devoted to audits.[8]

Although the audited individuals were selected at random as part of the project, the regular audit flag system was applied as in standard operating procedures. For each taxpayer, data contains information on income by line item as well as total taxes due before and after the audit. Also, the amount of income reported by third parties is available by line item, as well as flag indicators from the audit flag system.

In our case, we aggregate to a broad measure of income that we denote "total income", which is basically the sum of all income less deductions. Table 1 presents key descriptive statistics of the data. Average reported income is slightly less than 200,000 DKK[9] and on average about 63,000 DKK were due in taxes on reported income. On average, during the audits, reported income was adjusted upwards 1,664 DKK showing a clear asymmetry in reporting behavior "in favor" of the taxpayer. Adding the audit adjustment to reported income yields post-audit income, which is slightly less than the average amount of income reported by third-parties, corresponding to taxpayers on average having deductions that are not reported by third parties.

Although the average audit adjustment is very small relative to average post-audit income, it is large relative to average post-audit self-reported income.

[8]We do not use the sample of self-employed taxpayers in this study. In the entire experiment with all treatments and waves of audits a total of 25,020 wage earners and 17,764 self-employed were audited.

[9]In 2006, 1 USD was approximately worth 6 DKK.

Table 1. Tax Compliance of Danish Wage Earners, Income Year 2006.

Measure of income		Total income	Total taxes
Reported income	($\tilde{y}$)	193,277	63,178
		(1,906)	(841)
Audit adjustment (avg. over pos. and neg. changes)	($y - \tilde{y}$)	1,664	636
		(480)	(246)
Post-audit income	(y)	194,941	63,814
		(1,947)	(872)
Third-party reported income	(z)	195,618	
		(1,844)	
Self-reported income pre audit	(x)	-2,341	
		(584)	
Self-reported income post audit	(u)	-677	
		(711)	

Notes: The table shows descriptive statistics (means and standard errors) for a stratified random sample of 10,740 taxpayers denoted as wage earners (incl. unemployed, pensioners, etc.). Due to the stratification strategy employed by SKAT, the sample contains 74.6 percent "heavy" taxpayers (i.e., with high-complexity tax returns) and 25.4 percent "light" taxpayers, whereas the population has 32.6 percent heavy taxpayers and 67.4 percent light taxpayers.
Total income is defined as personal income + capital income − deductions + stock income + self-employment income + foreign income, where personal income consists of labor market income, social transfers, and pensions less labor market taxes and some pension contributions. Reported income is the sum of third-party reported income and self-reported income. Standard errors of means in parentheses. All estimates for wage earners are population weighted.
All amounts in DKK (1 USD ≈ 6 DKK in 2006).

3.3 Numerical Solution of the Model

The second order boundary value differential equation can be solved for any z using standard differential equation algorithms. The problem is solved backwards from the highest report, $\overline{x}$, to the lowest value of residual income, $\underline{u}$, and is initialized using $p(\overline{x}) = 0$ and $p'(\overline{x}) = \left(\frac{1}{1+\theta}\right) / (\overline{u} - \overline{x})$, where $\overline{x} \equiv x(\overline{u})$. However, since a positive mass of taxpayers are pooling their reports at $x = \underline{u}$, the solution must also satisfy the boundary condition that the tax agency's first order condition (FOC) is met in the pooling point at $x = \underline{u}$. We use a so-called "shooting method" algorithm to determine the λ that satisfies the boundary condition by, for each λ, calculating the deviation from 0 of the tax agency's FOC, and employing a minimization algorithm to find the λ that satisfies this condition almost exactly.

While mathematically and intuitively z is naturally understood to be a continuous variable described by the simultaneous distribution of u and z, we approximate the optimal allocation of the total audit budget on the domain of z by constructing a representative, evenly spaced grid. The equilibrium solution can be found by, for each z, solving for the within-audit-group evasion and audit strategy such that λ satisfies the tax agency's FOC for all z.

In practice, the solution method works well only in settings with many observations. The numerical solution algorithm is computationally intensive, especially in areas with very low density. That means we have to restrict income domains on which we solve the model, removing areas of near-zero density. For the purposes of simulating GMM moments, this also means we must restrict our sample to individuals within the restricted income domains on which the model is solved.[10]

3.4 Estimation and Results

In order to get point estimates of the deep parameters of the model – the share of honest taxpayers, Q, the penalty factor, θ, and the shadow-price on the audit budget,[11] λ – we match moments from the data predicted by the model to moments from the observed data using generalized method of moments (GMM).

Using the model described in Section 2, we predict evasion behavior of tax evaders in the sample based on the empirical bivariate distribution of third-party reported income and residual income,[12] $\hat{F}(z, u)$, and the parameters of the model. This allows us to compare moments of observed evasion, $u - x$, to the model's predictions. As with evasion behavior, these moments are complicated nonlinear functions of the parameters, so we solve the model numerically in order to predict evasion behavior and calculate moments.

To identify the parameters, we need to think carefully about which moments to use in the estimation. In particular, it turns out that identifying Q and θ together is a

[10] In Boserup and Pinje (2012) the model solution is found by calibration, varying one parameter only, which takes many fewer function evaluations. As the numerical solution is less computationally demanding, outlier removal need not be as restrictive, leaving a larger sample for the model to run on. In that sense the parameters in Boserup and Pinje (2012) and in the present paper are not directly comparable. See also Boserup and Pinje (2012) for a precise description of how the income domain on which the model is solved is defined.

[11] Given the income distribution, the size of the audit budget follows directly from the other parameters.

[12] We fit a mixed log-normal distribution with three component distributions. This approach is flexible enough to fit the data without producing "troughs" of zero density in the interior of the domain, which would cause the numerical solution method to break down.

challenging task. For example, in equilibrium, the share, Q, of honest taxpayers affects evasion of tax evaders positively, in short, by making it easier for tax evaders to hide among an increasing number of honest taxpayers, whereas the penalty rate, θ, works to discourage tax evasion by promising harsher penalties for detected evasion. Most other moments related to tax evasion have this property; if they are positively affected by the one parameter, they are negatively affected by the other, and identification of both parameters simultaneously fails.

However, interacting tax evasion with the share of evaders, $(1-Q)$, delivers a useful moment as a complement to the level of tax evasion for evaders. In the context of the model, $E[(1-Q)(u-x)|u-x>0]$ has the interpretation of mean evasion among all taxpayers, evaders and non-evaders, with at least some income not reported by third parties (i.e., $u \neq 0$). For the lack of a better word, we dub this moment *population level evasion*, although the population does not correspond to the universe of all taxpayers.[13] Whereas tax evasion for tax evaders is increasing in Q, population tax evasion is instead decreasing in Q as the mechanical effect of there being fewer tax evaders dominates the behavioral effect of fewer tax evaders evading more on average.[14] Meanwhile, θ discourages both mean evasion among evaders and at the population level. Thus, drawing on both the share of evaders and the level of evasion per evader allows for joint identification of Q and θ.

In addition to mean tax evasion among evaders and in the broader population, we need at least a third moment to identify parameters. We use the covariance for tax evaders between evasion and third-party reported income as the third moment. This moment is useful because it captures the degree to which high income income individuals evade more

[13]Of course, it is simple to calculate actual population evasion for all taxpayers. However, as the optimal weighting matrix used in the GMM procedure is based on matching individual contributions of the observations to each moment, it is infeasible to include taxpayers in the mass point at $u|z=0$. This is because the model we employ treats true residual income as an exogenous variable, meaning that if we were to introduce such taxpayers in the moment calculations, the systematic correlations in evasion behavior among evaders in the model and the data would be dominated by the noise of a stochastic allocation of taxpayers in or outside the mass point. In principle, this can be done, but it would take more data, both in the sense of observing many more evaders and in the sense of observing multiple correlates (that we do not have) to model the propensity of having zero residual income.

[14]Of course, this is only true as long as Q is fairly large, which is likely to be the relevant case in general, and certainly is the case in Denmark.

or less taxes relative to low income taxpayers. In turn, this reflects the relative intensity of audits among high income taxpayers versus low income taxpayers and renders this moment useful for distinguishing the effects of increasing the audit budget (lowering λ) and increasing θ.

These three moment conditions provide exact identification of the GMM estimator for (λ, θ, Q). The estimation is carried out as a two-step GMM procedure with optimal weighting matrix. The basic estimation results for the exact identified case are presented in column 1 of Table 2, displaying reasonable parameter values. Although the estimate of the penalty rate, θ, at 1.643 is not significant, the point estimate fits well with actual values applied in the Danish context, where the same factor is ½, 1, or 2 depending on the severity and the size of tax evasion found; see Boserup and Pinje (2012) for details. Likewise, the share of honest taxpayers, Q, is estimated at 0.889; in the data, the share of wage earners that do not under-declare income despite having some income that SKAT did not know of from third-parties is 0.855. Given the estimates of λ, θ, and Q, we can back out the implied budget value, B, which is 0.070, corresponding to 7.0 percent of all wage earners being selected for audit. This is higher than the total population audit rate of 4.2 percent reported in Kleven et al. (2011), but as taxpayers in our model are risk neutral, the budget size is likely compensating for this.

For robustness, we add a fourth moment, the covariance of reported residual income, x, and post-audit residual income, u, for tax evaders. The estimation results are presented in column 2 of Table 2. The parameter estimates are basically unchanged, but notice that θ is much better identified. Adding the extra moment improves identification, which is also reflected in the model passing the test for over-identifying restrictions. The over-identified model is our preferred specification.

3.5 Goodness of fit

The model predicts reporting behavior of taxpayers as well as the strategic behavior of the tax authorities. But in constructing goodness-of-fit measures, we rely solely on taxpayer behavior, because actual audit probabilities are unobserved.

Table 2. Estimating the parameters of the model using two-step GMM.

Parameter	(1)	(2)
λ	1.025	1.053
	(0.001)	(0.021)
θ	1.643	1.599
	(0.973)	(0.024)
Q	0.889	0.885
	(0.082)	(0.084)
Implied B	0.070	0.071
Over-identifying restrictions	0	1
Over-identification test (p-val.)	–	0.196
No. of obs.	797	797
	Goodness of fit $(R^2)^\dagger$	
Reported income	0.990	0.990
Reported residual income	0.752	0.752
Evasion	-0.120	-0.120

Notes: Robust standard errors in parentheses.
Data restrictions: total no. of obs. 10,740, unrestricted number of evaders 907, restricted number of evaders: 797, cf. Section 3.3.
Two-step GMM estimation was performed using an identity matrix as initial step weighting matrix and the optimal weighting matrix in the second step. Standard errors are robust to heteroscedasticity. The results presented are not sensitive to the choice of starting point of the estimation algorithm.
$^\dagger R^2$ is calculated as $1 - SS_{\text{err}}/SS_{\text{tot}}$.

Table 2 reports R^2 based on different income concepts. The model performs extremely well when we look at reported income, $\tilde{y} = z + x$, with an R^2 of 0.990. As the model predicts reporting behavior conditional on third-party reported income, i.e., x conditional on z, we calculate the R^2 based on the reported residual, x, alone. Again, the result is impressive at 0.752. To assess the fit of the model further, we find interestingly a negative R^2 for evasion – the variance in the prediction errors of evasion is greater than the variance in evasion itself.

From these results we may conclude that the model performs well in predicting reporting behavior mainly by incorporating third-party reporting. Also, the model explains the level of evasion with reasonable parameter values. Yet, on the individual level, the model fares less well. This is likely the result of our model simplistically only incorporating heterogeneity in terms of income portfolio (third-party reported versus residual income). The model does not include taxpayer heterogeneity in terms of the utility function (observables such as gender, age, etc.) or additional information in the audit probability function of the tax authorities (e.g., information drawn from the combination of line items used on a taxpayer's tax return). Yet, an attempt to use such information in this context is not feasible due to a relatively small number of observations.

4 Policy Experiments

To demonstrate the model's usefulness in simulating counterfactuals, we simulate an actual, ongoing policy change in Denmark in 2010-2012. SKAT implemented this change by mandating primarily banks and other financial entities to record buying and selling prices of each individual's assets and report this information to SKAT in a yearly manner.[15] Before the policy change, each taxpayer was required to self-report this information. For that reason, a relatively great deal of tax evasion took place on line items in this category and, as a consequence, this category of income was highly prioritized in SKAT's strategic disposition of resources from 2007 onwards.[16]

In comparing counterfactuals, we employ a measure of the average share of total income in the population that is subject to information reporting by third parties. We dub this the *information rate*, and it is defined by

$$I = 1 - \frac{\int |u| f(u) du}{\int |y| f(y) dy}, \tag{3}$$

which recognizes the fact that the tax agency can have third-party information on both positive and negative incomes. If all income is reported by third parties, I equals one. Conversely, if no income is reported by third parties, I is zero.

For 2006 data and actual policies, we find that the information rate for our sample of wage earners is 0.969, highlighting the massive amount of third-party reporting in the Danish tax system. We find that introducing third-party reporting on capital gains on stocks corresponds to an approximate 1.5 percentage points increase, giving an overall information rate of 0.985 for our sample of wage earners.

[15] Dividends were at the time already subject to third-party reporting from the same entities.

[16] See SKAT (2007, 2008, 2009, 2010, 2011), i.e., SKAT's official strategic dispositions of resources 2007-2011. Unfortunately, the documents are only available in Danish.

We can simulate the outcome of changing the reporting requirements by reassigning all self-reported stock income from the category of self-reported income to the category of third-party reported income, solving and simulating the model using conditional residual income distributions that reflect this policy change. In general, the change in the reporting requirement has the effect of decreasing the conditional variance of u given z, both by narrowing the domain of possible reports and by concentrating a larger measure of true residual income around 0.

We report the approximated changes in information rates as well as results of policy experiments in Table 3.

Table 3. Policy Experiment

Parameters	Baseline	Policy variant		
		Capital gains third-party reported	Revenue equivalent experiments	Parameter change (population equivalent)
Audit rate, B	0.071	0.071	0.131	0.060
Penalty rate, θ	1.599	1.599	>10	-
Honesty, Q	0.885	0.885	0.922	0.015
Information rate, I	0.969			
I, incl. capital gains		0.985		0.015
I, increased on general income			0.986	0.016
Population evasion, '000 DKK	0.230	0.153	0.153	

Notes: First column: Baseline parameters are the two-step GMM estimated parameters in the over-identified model. *Second column:* Simulation of the 2010-2012 policy change that implements third-party reporting for capital gains on stocks using baseline parameters. *Third column:* Counterfactual *ceteris paribus* parameter changes calibrated to have the same effect on population evasion as in column two. *Fourth column:* Parameter changes of the policy experiments in column three normalized to population scale. The change in B corresponds to the percentage points of extra taxpayers audited in the population of wage earners. Q is defined as the share of honest taxpayers among taxpayers with some income not reported by third parties. The added share of honest taxpayers in the population is therefore given by $\Delta Q \cdot (1 - M)$, where M is the share of taxpayers with their entire income reported by third parties. The population-equivalent change in I is interpretable as the share of average taxpayers whose income would be converted from fully self-reported to fully third-party reported.
Note that the simulated effects of the policy are estimated using the 2006 data described in Section 3.2. The actual policy change simulated in column 3 was implemented in 2010-2012.

In the first column we report the baseline (4 moment GMM) parameter estimates in terms of the audit rate, B, the penalty rate, θ, and Q and the baseline information rate 0.969.

In the second column we report the simulated effects of implementing third-party reporting on capital gains on stocks holding constant the parameters reported in the first column.[17] As mentioned above, we can see that the information increases by roughly 1.5 percentage points resulting in a 33 percent reduction in average population evasion from 230 DKK per capita to 153 DKK per capita.

In the third column we evaluate counterfactual policy experiments, varying in turn the audit budget, B, the penalty rate, θ, and the fraction of honest taxpayers, Q, in order that the reduction in population evasion be the same as the simulated reduction that results from the 2010-2012 policy change. In addition, we also calculate for comparison the necessary increase in the information rate if we, for each person, reduce proportionally residual income rather than specifically re-categorizing capital gains on stock income.[18]

First, we find that no "reasonable" value of the penalty rate can lower population evasion to the required level.[19] Penalties are ineffective because they potentially fall primarily on tax evaders that are subject to very low probabilities of audit. Note that this is consistent with results from the experimental literature (e.g., Alm, Jackson, and McKee, 1992a; Torgler, 2002), which generally fail to demonstrate a significant impact of penalties on tax evasion.

Next, we find that by increasing the audit budget, B, we can lower population evasion to the required level, but doing so requires almost doubling the fraction of audited taxpayers from approximately 7 percent to 13 percent.[20] The combination of the principle of incentive compatibility and a large mass of honest taxpayers reporting zero residual income constrains the tax agency from intensively auditing high-risk (i.e., high-income) taxpayers. Instead, the agency must preferentially increase the audit intensity among taxpayers reporting relatively low residual incomes and who are already subject to relatively high audit intensity.

When varying the honesty parameter, we find that the necessary value of Q is 0.922

[17] I.e., λ is allowed to vary such that the audit budget is held constant.

[18] I.e., for each person, i, change residual income by $(1-\delta)\cdot u_i$ such that counterfactual residual income is δu_i and counterfactual third-party reported income is $z_i + (1-\delta)\cdot u_i$.

[19] In fact, we have evaluated values of the penalty up to 10, compared to the Danish maximal rate of 2, and have failed to reach the required value. The equilibrium differential equation is extremely difficult to solve for large values of θ, so we do not pursue larger values.

[20] Recall from Section 2 that we normalize the per capita cost of audit to 1 such that the audit budget is equivalent to the equilibrium fraction of taxpayers subject to audit.

compared to the baseline estimate in which Q is equal to 0.885. Increasing the proportion of honest taxpayers is an effective policy variable because, like information collection, evasion is primarily reduced mechanically by the reclassification of income, in this case from income belonging to a potential tax evader to income belonging to an honest tax filer. This policy parameter has been studied extensively on the behavioral literature on tax evasion, suggesting fairness of the tax code, satisfaction with the provision of public goods, local democracy, whether tax authorities treat taxpayers respectfully, etc., can induce taxpayers to be voluntarily compliant (e.g., Pommerehne, Hart, and Frey, 1994; Feld and Frey, 2002; Frey and Feld, 2002). Interestingly, the Danish tax agency, SKAT, taking this literature to heart, actually underwent a strategic image change (roughly, branding itself as a public service rather than an adversary) in 2007 with the explicit goal of promoting tax compliance (SKAT, 2007). However, despite the behavioral literature on this topic, little research illuminates the actual degree to which tax agencies are able to increase voluntary compliance in this manner, especially in countries where honest reporting is as prevalent as in Denmark. E.g., Slemrod, Blumenthal, and Christian (2001) conduct an experiment, among other things, testing the effectiveness in deterring tax evasion of appeals to taxpayers' consciences. They are unable to find significant effects of the appeals on reporting behavior.

Lastly, the "general income" change in the information rate required to achieve the same level of population evasion is 0.986 compared to the rate of 0.985 when changing specifically capital gains on stocks from the residual income to the third-party reported income categories. The slightly larger necessary change in the information rate for the general-income change reflects the fact that tax evasion is especially rampant on line items related to capital gains on stocks, so mandating information reporting on these line items is somewhat more effective than a less focused increase in information reporting.

In the fourth and rightmost column we provide the parameter changes relevant to the second and third columns, translated into population equivalents (again, ignoring θ for which we have no point estimate for the necessary change and for which no such translation exists). For B and Q these changes are easily calculated as these parameters are already defined relative to the population of taxpayers. Therefore, the population equivalent change for B is simply 0.060, i.e., an additional 6 percent of the population must be audited in equilibrium to obtain the necessary change in population evasion. For Q we correct for the fraction of taxpayers with zero residual income to arrive at a necessary change in the fraction of compliant taxpayers of 0.015 or approximately 1.5 percentage points. Likewise, the change in the information rate can be re-interpreted as the additional number of entirely self-reported "average taxpayers" for whom we must have included their entire income in third-party reports.[21] Interestingly, the population equivalent changes for Q and I are roughly the same size, suggesting that the policy effect of changing reporting behavior for entire tax returns for relative few individuals (i.e., changing Q) is the same as changing reporting behavior on smaller parts of tax returns for relatively many individuals (i.e., changing information reporting).

Audit rates and penalty rates are policy instruments that work along the intensive margin of tax evasion. The simulations in Table 3 show that these instruments are far less effective in deterring tax evasion compared to third-party information reporting and the share of honest taxpayers in the population, which work along the extensive margin. E.g., an almost 100 percent increase in the audit budget is necessary to obtain an equivalent reduction in population evasion as increasing either the information rate or the fraction of compliant taxpayers by 1.5 percentage points.

The policy experiment suggests that information collection and the share of honest taxpayers are equally effective. The mechanism through which they work are very similar. Both entail a mechanical decrease in the mass of income available for potential tax evasion. However, the cost of the two are very dissimilar. The cost to the tax agency of increasing information collection are likely large, but the expenses are mostly one-off due to the, at least in Denmark, largely electronic nature of such reporting systems. In contrast, the potential tax revenue derived from increased information reporting is compounded from all future years. Although little is known about the possible avenues for increasing

[21]That is, "average taxpayers" in terms of how much residual income must be reclassified, for the average person, from residual to third-party reported income, as a fraction of the population, such that we obtain the necessary change in population evasion.

voluntary compliance among taxpayers, we can conjecture that the costs of doing so are variable, rather than fixed, and that the investment will not pay off (or, at least, deteriorate) in future years. Further, while increasing third-party information reporting can be implemented in a highly focused manner as in the 2010-2012 policy change, campaigns to promote honest reporting are likely less capable of reaching a select group of potential evaders.

Of course, the scope for information reporting is conditional on legislation allowing large scale individual-level information collection, which again is conditional on the degree of trust people have in government institutions not abusing such information. In some countries, e.g., Germany and the U.S., civil liberties legislation is much more prohibitive of the collection and utilization of individual level administrative data, making information reporting less attractive as a policy instrument, whereas, e.g., tax authorities in the Scandinavian countries make widespread use of third-party information reporting.

5 Conclusion

This paper provides a first structurally estimated model of tax evasion and enforcement. With estimated deep parameters within the realm of reason we are able to match moments of observed tax evasion.

Our results strongly suggest that policy instruments working along the extensive margin are superior to policy instruments working along the intensive margin in deterring tax evasion. The reason is that policies working through the intensive margin are not as effective at targeting the expected cost-benefit trade-off of taxpayers evading taxes by large amounts relative to that of the average taxpayer. In the existing literature, the penalty rate on tax evasion is ignored as a policy variable (an exception is Mookherjee and Png, 1989) under the assumption that penalties large enough to be effective are politically infeasible. Our results lend support to this standard practice as they suggest that the necessary penalty rate is very large. In addition, the apparent ineffectuality of policies working on the intensive margin is consistent with the experimental literature's findings that moderate changes in penalty rates and audit probabilities do not substantially affect fraudulent behavior.

Our results do not permit a conclusion as to whether increasing third-party information reporting or promoting honesty is superior in combating tax evasion. Basic intuition suggests that it is more straightforward to implement a more extensive system of third-party reporting. Tax agencies do not *ex ante* know which taxpayers are potential evaders, therefore they do not know to whom they should target a treatment for the purpose of promoting honest reporting. In addition, the experimental evidence of the effectiveness of using, e.g., moral suasion to deter tax evasion is mixed between negative results in field experiments (i.e., zero-effect results, e.g., for moral suasion in Slemrod et al., 2001; Torgler, 2004) and small to moderate effects (e.g., voting on penalties or shaming cheaters in Feld and Tyran, 2003; Coricelli, Joffily, Montmarquette, and Villeval, 2010, respectively) in laboratory experiments. Overall, these arguments come down in favor of information collection as most straightforwardly implementable and most effective policy to combat tax evasion.

The missing link in previous attempts to understand why taxpayers seemingly evade too little compared to average audit rates is information. Kleven et al. (2011) and Boserup and Pinje (2012) along with this paper show that taking account of the role of third-party information reporting to a large extent explains the observed levels of tax evasion. This paper also emphasizes that in order to explain the behavior of tax evaders, we need to incorporate the distribution and proportion of honest taxpayers. All in all, our paper suggests that there is a role for both the behavioral and the expected utility maximizing paradigms in the literature on tax evasion.

Bibliography

Allingham, MG and Sandmo, A, "Income tax evasion: A theoretical analysis." *Journal of Public Economics*, volume 1 (1972), pp. 323–338.

Alm, J, Jackson, B, and McKee, M, "Estimating the determinants of taxpayer compliance with experimental data." *National Tax Journal*, volume 45 (1992a), no. 1, pp. 107–114.

Alm, J, McClelland, GH, and Schulze, WD, "Why Do People Pay Taxes?" *Journal of Public Economics*, volume 48 (1992b), no. 1, pp. 21–38.

Andreoni, J, Erard, B, and Feinstein, JS, "Tax Compliance." *Journal of Economic Literature*, volume 36 (1998), no. 2, pp. 818–60.

Border, KC and Sobel, J, "Samurai Accountant: A Theory of Auditing and Plunder." *Review of Economic Studies*, volume 54 (1987), no. 4, pp. 525–40.

Boserup, SH and Pinje, JV, "Tax Evasion, Information Reporting, and the Regressive Bias Prediction." PhD dissertation, University of Copenhagen (2012).

Clotfelter, C, "Tax evasion and tax rates: An analysis of individual returns." *The Review of Economics and Statistics*, volume 65 (1983), no. 3, pp. 363–373.

Coricelli, G, Joffily, M, Montmarquette, C, and Villeval, M, "Cheating, emotions, and rationality: an experiment on tax evasion." *Experimental Economics*, volume 13 (2010), no. 2, pp. 226–247.

Erard, B and Feinstein, JS, "Honesty and Evasion in the Tax Compliance Game." *RAND Journal of Economics*, volume 25 (1994), no. 1, pp. 1–19.

Feinstein, J, "An econometric analysis of income tax evasion and its detection." *The RAND Journal of Economics*, volume 22 (1991), no. 1, pp. 14–35.

Feld, L and Tyran, J, "Tax evasion and voting: An experimental analysis." *Kyklos*, volume 55 (2003), no. 2, pp. 197–221.

Feld, LP and Frey, BS, "Trust Breeds Trust: How Taxpayers Are Treated." *Economics of Governance*, volume 3 (2002), no. 2, pp. 87–99.

Frey, B and Feld, L, "Deterrence and Morale in Taxation: An Empirical Analysis." CESifo Working Paper No. 760 (2002).

Kleven, HJ, Knudsen, MB, Kreiner, CT, Pedersen, S, and Saez, E, "Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark." *Econometrica*, volume 79 (2011), no. 3, pp. 651–692.

Mookherjee, D and Png, I, "Optimal Auditing, Insurance, and Redistribution." *Quarterly Journal of Economics*, volume 104 (1989), no. 2, pp. 399–415.

Phillips, MD, "Taxpayer Response to Targeted Audits." mimeo, University of Chicago (2010).

Pomeranz, D, "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax." mimeo (2010).

Pommerehne, W, Hart, A, and Frey, B, "Tax morale, tax evasion and the choice of policy instruments in different political systems." *Public Finance= Finances publiques*, volume 49 (1994), no. Supplement, pp. 52–69.

Reinganum, JF and Wilde, LL, "Income Tax Compliance in a Principal-Agent Framework." *Journal of Public Economics*, volume 26 (1985), no. 1, pp. 1–18.

SKAT, "Indsatsplan 2007." Report (2007).

———, "Indsatsplan 2008." Report (2008).

———, "Indsatsplan 2009." Report (2009).

———, "Indsatsplan 2010." Report (2010).

———, "Indsatsplan 2011." Report (2011).

Slemrod, J, "Cheating Ourselves: The Economics of Tax Evasion." *Journal of Economic Perspectives*, volume 21 (2007), no. 1, pp. 25–48.

Slemrod, J, Blumenthal, M, and Christian, C, "Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota." *Journal of Public Economics*, volume 79 (2001), no. 3, pp. 455–483.

Taxpayer Advocate Service, IRS, “National Taxpayer Advocate’s 2010 Annual Report to Congress.” (2010).

Torgler, B, “Speaking to theorists and searching for facts: Tax morale and tax compliance in experiments.” *Journal of Economic Surveys*, volume 16 (2002), pp. 657–683.

———, “Moral suasion: An alternative tax policy strategy? Evidence from a controlled field experiment in Switzerland.” *Economics of Governance*, volume 5 (2004), no. 3, pp. 235–253.

Chapter 3

Intergenerational Wealth Mobility: Evidence from Danish Wealth Records of Three Generations*

Simon Halphen Boserup
University of Copenhagen

Wojciech Kopczuk
Columbia University, CEPR, and NBER

Claus Thustrup Kreiner
University of Copenhagen and CEPR

March 2013

Abstract

We provide empirical evidence on the intergenerational mobility of wealth using administrative wealth records of three generations of Danes. Our preferred estimate for the intergenerational wealth elasticity (IWE) is 0.2 (and 0.27 when limiting attention to those with positive wealth only). We construct a theoretical framework that allows for understanding the variability of the IWE across time, samples, and countries. Our framework highlights that the IWE can be interpreted as the weighted average of elasticities corresponding to different sources of intergenerational correlation that may in principle vary in importance across different contexts. However, we find that the IWE is surprisingly stable when estimated for different age groups, when using parents-grandparents pairs instead of children and parents, when eliminating bequests, and when explicitly shutting down many of the potential channels behind intergenerational wealth mobility, including income and education. This suggests that parental wealth is a sufficient statistic for the channels that we control for and those that vary across different samples, that is, the effect of these parental characteristics on wealth of children can be summarized by their effect on wealth of parents. By exploiting information for three generations we find that the standard child-parents elasticity severely underestimates the long-term persistence in the formation of wealth across generations. We show that either the true elasticity is significantly underestimated or that grandparental characteristics matter beyond information incorporated in parental characteristics. We also find evidence supporting the presence of a persistent dynastic component.

*We thank Søren Leth-Petersen for helpful discussion. Benjamin Ly Serena provided excellent research assistance. Financial support from WEST is gratefully acknowledged. Contact information: simon.h.boserup@econ.ku.dk, wojciech.kopczuk@columbia.edu, ctk@econ.ku.dk.

1 Introduction

The main objective of the literature studying intergenerational mobility is to analyze the extent to which economic well-being is related across generations, and to try to understand the underlying mechanisms (Piketty, 2000). In other words, the interest is both in the strength of the relationship between lifetime economic resources/consumption possibilities of children and parents, and in determining whether the statistical association is governed by mechanical transmission of abilities and traits from parents to children or by active decision-making of parents, such as investment in human capital of their children. Intergenerational correlations in earnings and income have been extensively studied (see e.g. the surveys of Solon (1999) and Black and Devereux (2011)), but much less evidence exists on intergenerational wealth mobility, even though wealth measured at some point in life may be a proxy for lifetime economic resources that is as good as income (or even better). An important exception is Charles and Hurst (2003), hereafter C&H, who use wealth data from the Panel Study of Income Dynamics (PSID) to estimate the elasticity of child wealth with respect to parental wealth for the United States. C&H obtain an age-adjusted intergenerational wealth elasticity (IWE) of 0.37 before transfer of bequest (i.e., both parents are alive), which is at the low end of the range of estimates of the intergenerational income elasticity for the US, and find that around 2/3 of the elasticity is accounted for by income and asset ownership.[1]

In this paper, we provide new empirical evidence on the intergenerational mobility of wealth using Danish administrative wealth records of three generations observed for the entire Danish population. Our data has several advantages. First, in our basic sample we have more than 1 million child-parents pairs compared to around 1500 in C&H. This allows us to include a rich number of controls and to split the sample into age-cohorts while still obtaining very precise estimates. Second, the use of administrative data removes problems of attrition and measurement errors that often plague survey studies.[2] Third, we are able to link comprehensive data for three generations (child-parents-grandparents). The information about an additional generation allows us to test

[1] Charles and Hurst (2003) review a few older studies looking at the intergenerational correlation of wealth. These studies have looked at small non-representative samples with few observations and poor data quality.

[2] Recent research has documented large survey measurement errors in key economic variables such as income and that the errors are correlated with conventional covariates implying that errors are non-classical (Kreiner et al., 2012).

whether child-parents relationships are stable and may also be used to address whether the simple IWE estimate provides an accurate measure of the degree of persistence in the wealth process across generations.

We contribute to the literature in a number of ways. We start by providing a theoretical framework capturing the relevant mechanisms that may create correlation of wealth across generations and that lays the foundation for the empirical analysis. The standard Becker and Tomes (1979, 1986) framework focuses on links in economic outcomes across generations due to correlation of abilities across generations and due to parental investment in human capital of children. These channels determine the association of income across generations and therefore are also likely to influence the relationship between wealth across generations. In addition to these mechanisms, however, wealth may be correlated across generations because of direct transfer of wealth from previous generations (inter vivos or through inheritance), because of correlations in patience and risk preferences creating differences in saving propensities (Stiglitz, 1969), or because of correlations in investment ability and the corresponding return (for example, due to stock market participation, entrepreneurship, or attitudes toward borrowing). Individual wealth observed at a given point in time therefore reflects in a complicated way economic resources, abilities, and traits inherited from the previous generation, and we show theoretically how the IWE coefficient is related to these underlying mechanisms.

The first part of our empirical analysis follows C&H. We find an (unconditional) child-parent age-adjusted wealth elasticity equal to 0.27 when focusing on indviduals with positive wealth as they do. This estimate is considerably lower than the IWE estimate of C&H. When we address the sample selection bias by using the full population, i.e., not removing child-parents pairs with negative wealth of either child or parents (this is our baseline specification), the estimate is even lower at 0.19.[3] Recent studies for France (Arrondel, 2009) and Denmark (Kolodziejczyk, 2011) also find an IWE estimate of 0.2. The lower IWE estimate for Denmark is compatible with cross-country studies of intergenerational earnings elasticities that find low elasticities in the Nordic countries

[3] C&H use a standard log transformation of the data and therefore remove non-positive values. However, negative wealth may be optimal from an economic theory point of view in certain circumstances, in particular for young individuals, and many individuals in our data have negative wealth. To allow for negative wealth, we apply the inverse hyperbolic sine transformation (IHS) of the data, which yields identical results as the log transformation for both positive wealth and the absolute value of negative wealth taken in isolation, but in addition allows for parsimoniously combining the two.

and the lowest elasticity for Denmark equal to 0.12 (Björklund and Jäntti, 2009).

Life-cycle variation has proven to be an important consideration in the measurement of intergenerational mobility in income that appears to be age-dependent (Haider and Solon, 2006). Similarly, there may be substantial variation in the IWE depending on the age of the child and the parents. To address this issue we run separate regressions for each age-cohort of the children and for each age-cohort of the parents measured at the year when the child was born, respectively. The IWE coefficients are very precisely estimated and all lie within the range 0.16–0.22, revealing a surprisingly stable relationship between child wealth and parental wealth. In our baseline estimates, therefore, we can abstract from age-dependent differences in wealth correlation and proceed by pooling all cohorts while flexibly controlling for age of parents and children to adjust for the life-cycle patterns.

We obtain a similar conclusion when introducing additional covariates and when splitting the sample. When we add covariates that themselves have a strong explanatory power, such as number of siblings dummies, income level, education length, and portfolio composition dummies, the child-parents wealth elasticity falls but not by a lot. For example, income and education of children and parents can only explain 8 percent of the original IWE estimate. The strongest effect comes from financial composition dummies that may explain 25 percent of the IWE, indicating that correlation in household finance behavior across generations may be important, but all explanatory variables taken together can only explain up to 30 percent of the IWE. We also split the sample according to parental age. This has quite a large impact on the estimates of the effect of child and parental income on child wealth, but only a small effect on the IWE estimate.

Our theoretical framework suggests that robustness of our estimates to inclusion of income and education (and to a lesser extent of financial composition dummies) is consistent with parental wealth being a sufficient statistic for the mechanisms behind intergenerational correlation of wealth that these variables proxy for.

In the second part of our empirical analysis, we exploit availability of information about grandparents. We start by addressing whether the relationship between child wealth and parental wealth is stable over time/generations. Recent research has documented substantial changes over the long run in top income shares, in the relative importance of capital and labor income at the top of the income distribution, and in the evolution of inheritance (Atkinson, Piketty and Saez, 2011; Piketty, 2011) and, similarly, there may be long run forces that reduce or increase the strength of the relationship between wealth of children and wealth of parents. In light of our theoretical framework, this would be so when parental wealth is not a sufficient statistic for parental influence, so that changing importance of different mechanisms behind intergenerational correlation of wealth translates into changing the IWE. When estimating the intergenerational elasticity of parental wealth with respect to grandparental wealth we obtain an estimate of the IWE that is very similar to the estimate obtained from children-parents pairs. This is even more striking when taking into consideration that the generations differ in many other respects than just age, and indicates that the cumulative importance of the underlying mechanisms governing the intergenerational relationship in wealth is quite stable.

Next, we look at the correlation between children and grandparents in isolation. If parental wealth is a sufficient statistic for all previous generations then the coefficient on grandparental wealth should be the child-parents IWE raised to the power of two. We obtain a coefficient that is more than three times as high and very precisely estimated. This indicates that the degree of persistence across generations is higher than what is reflected in the IWE estimate. A reason may be that the underlying processes relating wealth across generations have longer memory than just one generation, for example because grandparents have a direct impact on their grandchildren (Solon, 2012). When including both parental and grandparental wealth in the regression, we obtain a significant and sizable coefficient on grandparental wealth, while the child-parents coefficient only falls a little. This is consistent with the hypothesis of more persistence and the underlying processes relating wealth across generations having longer memory than just one generation.

Another likely reason for the significant coefficient on grandparents is measurement errors in wealth or, in the same vein, that wealth is a noisy signal of abilities and traits transmitted across generations and of the size of money transfers given from generation to generation. Attenuation bias created by transitory components in the economic outcomes has been a major concern in the intergenerational income mobility literature since the influential contribution of Solon (1992), and it is common to take averages over some years as we have done to reduce the importance of transitory components. However, this method may only remove a small part of the attenuation bias if the transitory component has some persistence (Mazumder, 2005) and would not work at all to remove a bias from random "fixed effects", e.g. you are born lucky in terms of ability given your family background without transmitting this to your children. In the theory section, we demonstrate how idiosyncratic variation in wealth may bias the IWE estimates downward, but also discuss conditions under which it is appropriate to use wealth of grandparents as an instrument for parental wealth in order to obtain consistent estimates. When we redo the first part of the empirical analysis, we consistently obtain estimates of the IWE in the range 0.6-0.7, which is more than three times as high as the original ordinary least squares estimates. This points again to a much higher degree of persistence in the wealth formation across generations.

The exclusion restriction that allows for using grandparental wealth as an instrument for parental wealth is strong. Hence, the results indicate one of the two possibilities: either IWE is significantly underestimated when using the conventional approach or the effects extend for more than just a single generation. In either case, we interpret it as indicating that the extent of intergenerational mobility is substantially lower than a naïve estimate would indicate.

As an alternative approach, we further inquire into the degree of persistence by allowing for dynasty fixed effects. The standard Becker and Tomes framework, underlying intergenerational empirical analyses, assumes that economic outcomes of all generations converge towards the same constant. But as demonstrated by Stiglitz (1969) and in our theory section, people may belong to different classes/dynasties who differ with respect to productivity or savings behavior, implying that wealth, without any new shocks, converges towards a steady state distribution rather than a constant. Our empirical analysis allowing for fixed effects indicates that the model with a common constant is misspecified and therefore provides further evidence to the conclusion that the conventional IWE estimates underestimate the degree of persistence.

The remaining part of the paper is organized as follows. Section 2 provides a theoretical framework to understand the relevant mechanisms underlying the correlation of wealth across generations and the challenges in estimating the degree of persistence across generations. Section 3 describes the construction of the data sets and provides summary statistics of key variables. Section 4 describes the results of the empirical analysis. Finally, Section 5 offers concluding remarks and an appendix provides additional details concerning the data.

2 A theory of the correlation of wealth across generations

We start by providing a simple conceptual framework for understanding the relationship between wealth of different generations. Let's denote w_g to be wealth of generation g. A member of family i that belongs to generation g maximizes lifetime utility $u(\{C_a^g, q_a^g\}_{a=o}^{a=A}; \varepsilon_{ig})$, where C_a^g is consumption at age a, where $0 \leq a \leq A$, q_a are transfers to the subsequent generation made at age a (inter vivos when $a < A$ and bequests for $a = A$), and ε_{ig} is the set of taste parameters characterizing preferences of this individual. Optimization is subject to the initial wealth of $W_0^g = 0$ and the set of budget constraints for $a > 0$

$$W_a^g = W_{a-1}^g \cdot (1 + r_{ig} + \gamma_a) + y_a^g - C_a^g + B_a^g - q_a^g,$$

where W_a is wealth at age of a, y_a^g is income at the age of a, B_a^g is the bequest received from the previous generation at age of a (with one member per generation, $B_a^g = q_{\tilde{a}}^{g-1}$ where $\tilde{a} - a$ is the age difference between the parent and the child), r_{ig} is the mean rate of return specific to that particular individual (to allow for varying investment strategies), and γ_a is the mean-zero deviation from the normal rate of return. Consequently, the level of wealth can be expressed in terms of the exogenous parameters of the problem as

$$W_a^g \left(\{q_b^{g-1}\}_{b=0}^A, \{y_b^g\}_{b=0}^A, r_{ig}, \varepsilon_{ig}, \{\gamma_b\}_{b=0}^a\right). \tag{1}$$

In other words, wealth at any given age depends on the history and future (expected) transfers from parents, the history and future own income, the rate of return, preference parameters, and stochastic shocks.

We are interested in understanding the relationship between wealth of members of different generations. Most simply, one can observe the statistical association that may be, for example, described using the correlation coefficient $\mathrm{corr}(W_a^g, W_{\tilde{a}}^{g-1})$, where a and $\tilde{a}$ are ages at which we observe wealth of children and parents, respectively, or the empirical association between some transformation of observed wealth levels $\frac{dw_a^g}{dw_{\tilde{a}}^{g-1}}$, where (if we focus on positive wealth) we could use for example $w = \ln(W)$ or (as we will do in what follows) we can use the inverse hyperbolic sine transformation $w = \log(W + \sqrt{W^2+1})$ that behaves as $\pm\log(|W|)$ everywhere with the exception of in the neighborhood of zero.

Let's focus attention on

$$\frac{d\,w_a^g}{d\,w_{\tilde{a}}^{g-1}},$$

which we will refer to as the intergenerational wealth elasticity (IWE). Wealth of parents does not enter directly as an argument of child's wealth (Equation (1)), but rather it may be either correlated or causally related to the other arguments. Hence, it should be intuitive that this is not a well-specified concept, without taking a stand on why $w_{\tilde{a}}^{g-1}$ varies. By investigating the determinants of wealth in Equation (1), one can identify a number of reasons why wealth could co-vary:

- transfers from parents q_b^{g-1} are a function of their own characteristics and hence correlated with parental wealth
- incomes of parents and children may be correlated for many reasons extensively analyzed in the literature
- rates of return and preferences may be correlated across generations

Depending on the relative importance of these factors, we may expect to see different relationships at different times and places and different sample choices unless they happen to correspond to precisely the same strength of correlation in wealth. We expand on and formalize this intuition in what follows.

In order to make progress and introduce our empirical specification, we simplify our framework to resemble what has been sometimes dubbed the "mechanical" approach in the literature (Goldberger, 1989). It amounts to positing a statistical relationship between variables of different generations without explicitly specifying the nature of the individual optimization problem. We adopt this terminology but note that the framework we described is an optimization framework, except that we have refrained so far from imposing additional assumptions so that the framework serves only to identify the relevant variables. Our main simplification in what follows is approximation of the (potentially nonlinear) relationship rather than assuming away optimization.

In order to operationalize our empirical approach, we linearize our original wealth formula as

$$w_a^g \approx \alpha + \beta_q q_a^{g-1} + \beta_y y_a^g + \beta_r(1 + r_{ig}) + \beta_\varepsilon \varepsilon_{ig} + \eta_a^g, \tag{2}$$

where η is an error term incorporating rate of return shocks and approximation errors that is assumed (very strongly) orthogonal to the other variables and we posit that the current values of right hand side variables are sufficient to summarize the relationship with wealth. While we write all terms in a linear fashion to simplify notation, they can correspond to transformations (e.g., log or IHS) of the original variables.

When we estimate the statistical relationship between w_a^g and w_a^{g-1} as

$$w_a^g = \beta_w w_a^{g-1} + \omega, \tag{3}$$

the coefficient β_w is going to reflect the average impact of w_a^{g-1} running through all four possible channels: bequests, income, preferences, and rate of return. Our objective in the rest of this section is to clarify the interpretation and estimation of β_w.

In what follows, we abstract from the age aspect of inter-generational relationship (we will return to considering it in the empirical work). It will prove useful to write the determinants of wealth as a vector $x_g = (q^{g-1}, y^g, 1 + r_{ig}, \varepsilon_{ig}, \eta_g)$, with wealth expressed as

$$w_g = x_g \beta + \zeta_g,$$

where ζ_g is white noise measurement error. Finally, we specify the law of motion for determinants of wealth as

$$x_g = x_{g-1}\Xi + \nu_g\,. \tag{4}$$

We assume for now that it is autoregressive of order one, but deviations from this assumption will be important to consider in what follows. Equation (4) describes the relationship between characteristics of subsequent generations other than wealth. In what follows we will consider a special case when the relationship between x_{g-1} and x_g runs through wealth — that special case imposes structure on matrix Ξ but otherwise yields the law of motion (4).

Note that we have introduced many potentially unobservable sources of variation in wealth. First, taste parameters ε are likely unobservable directly (though perhaps they may be proxied for). It is natural to consider and test whether correlation in preferences is a source of intergenerational persistence. Second, we split the rate of return into two separate components. One is the normal rate of return — this is akin to a preference parameter in the sense of reflecting traits of individuals that lead them to select investments with varying outcome, although it is conceptually distinct since it represents manifestation of preferences through choices rather than preferences themselves. The other one is η that corresponds to deviations from the normal rate of returns (and approximation errors that we abstract from). We assume that it is an idiosyncratic components so that η is assumed not to be correlated across generations ($\Xi_{\eta\eta} = 0$). We include it though in x_g, because it is certainly possible that such random shocks to wealth do in fact have an impact on the subsequent generations by influencing other variables. The final source of randomness is ζ — this is assumed to be measurement error or other sources of variation in wealth that have no consequence for the subsequent generations.

Using this notation (and expressing all variables in terms of their deviation from the mean in order to eliminate the constant term), we can re-write wealth as follows:

$$w_g = x_g\beta + \zeta_g = x_{g-1}\Xi\beta + \nu_g\beta + \zeta_g = \beta_w x_{g-1}\beta + \xi_g = \beta_w w_{g-1} - \beta_w \zeta_{g-1} + \xi_g \qquad (5)$$

where β_w is a linear projection (linear regression coefficient) of $x_{g-1}\Xi\beta + \nu_g\beta + \zeta_g$ on $x_{g-1}\beta$ and ξ_g is the corresponding residual. In other words, β_w summarizes the statistical relationship between wealth of the two generations that runs through the set of characteristics in x (with orthogonal noise terms excluded).

More explicitly, using the standard OLS formula $(X'X)^{-1}X'Y$ with $X = x_{g-1}\beta$ and $Y = x_{g-1}\Xi\beta + \nu_g\beta + \zeta_g$ and noting that $\nu_g\beta + \zeta$ is orthogonal to $x_{g-1}\beta$, we can write the expected value of β_w as

$$\begin{aligned} \mathrm{E}[\beta_w \mid x_g, x_{g-1}] &= \left((x_{g-1}\beta)' x_{g-1}\beta\right)^{-1} \cdot (x_{g-1}\beta)' \cdot (x_{g-1}\Xi\beta) \\ &= (\beta' \cdot (x'_{g-1}x_{g-1}) \cdot \beta)^{-1} \cdot \beta' \cdot (x'_{g-1}x_{g-1}) \cdot \Xi\beta, \qquad (6) \end{aligned}$$

showing how β_w summarizes the complicated web of influences that links wealth across generations. It reflects the causal relationship between relevant arguments of Equation (2) (Ξ) that can be mapped to one-dimensional wealth within each generation via β. In general, the mapping of wealth across generations depends on the distribution of characteristics in the population x. This is because there are many channels through which characteristics of subsequent generations relate, and the *population* relationship between wealth of different generations reflects different components of Ξ depending on variation of x_{g-1} in the population. This is reflected by the presence of $x'_{g-1}x_{g-1}$ in the equation.

The coefficient β_w is as close as one can get to a "structural" intergenerational correlation of wealth if one insists on summarizing it by a single parameter. In the special case where x_{g-1} is one dimensional, so that Ξ is a scalar, it is easy to show that $\beta_w = \Xi$, implying that the β_w-coefficient reveals the (single) structural parameter.

2.1 Interpretation of β_w

More generally, β_w is a relationship that is specific to a given population and it does not have a direct causal interpretation. This is because the source of variation in wealth (w_{g-1}) matters: depending on which component of x_{g-1} is responsible, the effect will vary accordingly. β_w does though reflect the effect of variation in wealth in the population that combines influences from a variety of sources.

To illustrate the logic of this formula, consider the following example

Example 1 *Suppose that elements of x are uncorrelated with each other ($x_g' x_g$ is diagonal), Ξ is diagonal and only two elements of β, indexed by i and j are non-zero. Then straightforward manipulation yields*

$$\beta_w = \frac{\beta_i^2 \sigma_{ii}^2 \Xi_{ii} + \beta_j^2 \sigma_{jj}^2 \Xi_{jj}}{\beta_i^2 \sigma_{ii}^2 + \beta_j^2 \sigma_{jj}^2},$$

where σ_{ii}^2 is the (i,i) element of $x_g' x_g$ (variance of ith element of x_g, denoted by x_{gi}) and Ξ_{ii} is the (i,i) element of Ξ.

In this particular example, the correlation of wealth has its source in two different channels, i and j. A marginal increase in $x_{g-1,i}$, translates into a β_i increase in wealth of generation $g-1$ (w_{g-1}) and into a $\beta_i \Xi_{ii}$ increase in wealth of generation w_g. Hence, the intergenerational elasticity driven by this source of variation is Ξ_{ii}. Analogously, the intergenerational elasticity driven by the jth element is Ξ_{jj}. When both of the sources are present at the same time, β_w has to reflect both of these sources and their contribution depend on the relative magnitudes of $\beta_k^2 \sigma_{kk}^2$ ($k = i, j$). This is intuitive: what matters is the extent to which a given channel is responsible for variation in wealth — this is affected by the extent to which the given channel influences wealth (β_k) and the extent of variation in the underlying characteristics, σ_{kk}.

Note that even when one is willing to assume that β and Ξ are structural parameters, there is no single structural intergenerational elasticity here. Different societies at different points in time may differ with respect to the extent of variation in determinants of wealth — variation in bequests, tastes, education, rate of returns may all vary over time with institutions, policies, culture, etc. Each of such situations will correspond to a different weighted average of Ξ_{ii} and Ξ_{jj} and as will the measured IWE.

Another way of phrasing it is that the intergenerational wealth elasticity is the average treatment effect corresponding to changes in wealth of the prior generation. Since prior wealth may be varying for a multitude of reasons (corresponding to different impacts of the treatment — change in prior wealth), the corresponding estimate will reflect a weighed average impact with the weights being sample-dependent.

The example is of course restrictive, but illustrative of the logic of Equation (6) that characterizes the general case.

In the special case when $\Xi_{ii} = \Xi_{jj}$ the source of variation does not matter because the intergenerational elasticity stemming from the two different sources happens to be the same. A priori, this does not seem likely, but we next discuss a more general case where independence of the source of variation can stem from a less arbitrary assumption.

2.2 Parental wealth as a sufficient statistic

As discussed above, in general there is no single intergenerational elasticity of wealth because different sources of variation in parental characteristics may in general translate into different strengths of intergenerational association. However, it is possible to identify a special case when it is not so. Imagine that parental wealth is a sufficient statistic for the effect of parental characteristics on wealth of children. That is, retaining the linear structure, suppose that

$$x_g = (w_{g-1} - \zeta_{g-1}) \cdot \Gamma = w_{g-1}\Gamma - \zeta_{g-1}\Gamma$$

for some vector Γ (and note that we are using observed wealth net of ζ — the measurement error term). Then, $x_g = x_{g-1}\beta\Gamma - \zeta\Gamma$, so that — using our general notation — $\Xi = \beta\Gamma$ and $\nu = \zeta\Gamma$. Substituting into Equation (6) yields

$$\mathrm{E}[\beta_w \mid x_g, x_{g-1}] = (\beta' \cdot (x'_{g-1}x_{g-1}) \cdot \beta)^{-1} \cdot \beta' \cdot (x'_{g-1}x_{g-1}) \cdot \beta \cdot \Gamma \cdot \beta = \Gamma\beta,$$

which can also be seen directly from $w_g = x_g\beta + \zeta_g = w_{g-1}\Gamma\beta + \zeta_g - \zeta_{g-1}\Gamma$.

When parental wealth is a sufficient statistic for the characteristics of parents, the IWE is equal to $\Gamma\beta$ regardless of the source of variation in wealth — the formula for

$\mathrm{E}[\beta_w]$ is independent of x_{g-1}.

2.3 Testing the importance of different channels

So far, we have clarified the theoretical relationship between wealth of different generations while allowing for many different channels of interactions. Consider now the possibility of directly observing some of them. Specifically, suppose that we partition $x = (x^1, x^2)$ with x^1 unobservable and x^2 observed. Similarly, let's denote by Ξ_{11} and Ξ_{12} the partitions of matrix Ξ that determine x^1: $x_g^1 = x_{g-1}^1 \Xi_{11} + x_{g-1}^2 \Xi_{12} + \nu_g^1$. Then,

$$\begin{aligned} w_g &= x_g\beta + \zeta_g = x_g^1\beta + x_g^2\beta + \zeta_g = x_{g-1}^1\Xi_{11}\beta + x_{g-1}^2\Xi_{12}\beta + x_g^2\beta + \nu_g^1\beta + \zeta_g \\ &= \beta_w^1 x_{g-1}\beta + x_{g-1}^2\beta_{w,g-1}^2 + x_g^2\beta_{w,g}^2 + \xi^1 \\ &= \beta_w^1 w_{g-1} + x_{g-1}^2\beta_{w,g-1}^2 + x_g^2\beta_{w,g}^2 - \beta_w^1\zeta_{g-1} + \xi^1, \end{aligned} \tag{7}$$

where β_w^1 is a linear projection of $x_{g-1}^1\Xi_{11}\beta$ on $x_{g-1}\beta$ while partialling out the effect of x_{g-1}^2 and x_g^2 (with $\beta_{w,g-1}^2$ and $\beta_{w,g}^2$ being the resulting coefficients). In other words, by controlling for some subset of characteristics of both parents and children (note that controlling for both at the same time is important), we can zoom in on the effect of the remaining characteristics that is orthogonal to the observed ones.

2.4 Measurement error

Simple inspection of relationships (5) or (7) reveals that OLS of w_g on w_{g-1} will not in general estimate β_w (or β_w^1). This is because the error term ζ_{g-1} will be correlated with w_{g-1} — recall that $w_{g-1} = \beta \cdot x_{g-1} + \zeta_{g-1}$ or, more explicitly, the relationship is described by Equation (2). The attenuation bias due to the presence of ζ will not be a problem only when $\mathrm{var}(\zeta) = 0$. One may of course make assumptions to that effect, but absent these assumptions, our estimate of β_w will be biased.

Assuming that the measurement error term ζ_{g-1} is not serially correlated, characteristics of preceding generation x_{g-2} can serve as an instrument for w_{g-1}. In particular, a natural instrument to consider is w_{g-2}. Under the simple autoregressive structure of order one that we assumed, the exclusion restriction is satisfied and this is a valid instrument.

Similarly, specific components of x_{g-2} could be used as instruments. Note though that different instruments will identify different average treatment effects, i.e., in our case they will identify the weighted average of effects going through the channels that the instrument covaries with.

2.5 Multiple generations

We also have assumed so far that the characteristics follow the simple autoregressive structure of order one. Two natural generalizations are to consider a higher order autoregressive structure and to consider a fixed (or very persistent) family component.

Suppose first that x_g follows a second-order autoregressive process, $x_g = x_{g-1}\Xi^1 + x_{g-2}\Xi^2 + \nu_g$, while we continue to have $w_g = x_g\beta + \zeta_g$. We can define $y_g = (x_g, x_{g-1})$, $y_{g-1} = (x_{g-1}, x_{g-2})$, $\Xi^y = [\Xi^1, \Xi^2]$ so that $y_g = y_{g-1}\Xi^y + \nu_g$ and $w_g = y_g\beta^y + \zeta_g$ where $\beta^y = (\beta, 0, \cdots, 0)$. Substituting y_g, Ξ^y, β^y throughout in place of x_g, Ξ and β allows for repeating the whole argument, but now with the set of characteristics expanded to include grandparental ones. On the conceptual level, grandparental characteristics become yet another determinant of IWE.

In the presence of a family fixed effect, we can analogously expand the definition of x_g to include the constant family term.

While both of these extensions are straightforward for the purpose of understanding the IWE β_w, they raise additional estimation issues. First, they invalidate the possibility of using grandparental characteristics as an instrument. Second, eliminating the fixed family component in order to understand the importance of this channel is complicated due to the presence of lagged dependent variable structure.

2.6 Empirical strategy

We will proceed as follows. First, we will estimate *population* β_w. We will follow it by considering various subsamples and different generations in order to test the robustness of this estimate. Sensitivity of the results to different sample choices would indicate that the relative importance of different channels varies depending on the subsample. Lack of sensitivity would indicate either that wealth is a sufficient statistic or that, in fact, different samples considered happen to correspond to a similar mix of channels behind intergenerational correlation of wealth.

Next, we will consider controlling for different channels by including the corresponding characteristics of parents and children, and investigating the sensitivity of IWE to these choices. As we discussed, this approach allows for shedding light on whether the effect of the particular channel can be summarized through its impact on wealth.

Finally, we will consider grandparental characteristics to investigate the relevance of interactions that run across multiple generations.

3 Data

Our empirical analysis is based on data from several public administrative registers gathered by Statistics Denmark and linked together using personal identification numbers. Every citizen in Denmark is assigned a unique personal identification number at birth and the identification numbers of the mother and the father are registered for all Danes born in 1960 and onwards.[4] This enables us to combine different data sources at the individual level and to link data across generations.

The data on individual wealth and income is based on administrative tax return records, rather than survey questions as in Charles and Hurst (2003) and most other studies on intergenerational mobility. The Danish Tax Agency (SKAT) collects, in addition to information of various income sources, information about the values of asset holdings and liabilities measured at the last day of the year for all Danes, and the bulk of the wealth components are third-party reported.[5] The available pieces of information at Statistics Denmark are the aggregate value of assets and liabilities, respectively, covering the period 1980 to 2011, and from 1997 and onwards it is also possible to obtain complete portfolio information with respect to the value of bonds, stocks, cash in banks, house, mortgage loans and sum of other loans. Another attractive feature of the wealth data is that the information is not top coded.

The information about the value of financial assets and liabilities at the end of the year is reported to the tax authorities by banks, other financial institutions, and some government institutions, while the cash value of property is assessed by the tax authorities, based on detailed information of the property, and used for taxation of the imputed rent on the property. The third-party reported value of assets includes all deposits, stocks, bonds, value of property, and deposited mortgages. Pension funds are not included. The

[4] Registrations of parents exist before 1960 but are incomplete.

[5] The tax authorities use the income information to generate pre-populated tax returns. The information on wealth was originally used to compute the wealth tax, whereas today it is used by the tax agency to cross check if the reported income level is consistent with the change in net-wealth during the year under the assumption of a given estimated consumption level. A recent study by Kleven et al. (2011) reveals, using a large scale randomized tax auditing experiment constructed in collaboration with the Danish tax authorities, only small differences between the third-party reported income items and the corresponding items on the final tax return. This indicates that the third party reported information of the Danish Tax Agency is of a very high quality.

third-party reported value of liabilities includes debt in financial institutions, mortgage credit debt, credit and debit card debt, deposited mortgage debt, student debt and debt in The Mortgage Bank (a public institution), debt to financial corporations, debt to the Danish municipalities and other liabilities such as unpaid taxes and mortgage debt, which are not deposited.

From 1980 to 1996, Denmark had a wealth tax, and taxpayers had to self-report car values, boat values, caravan values, title deed of cooperative dwellings, premium bonds, cash deposits, stocks (both listed and non-listed thereby including privately held companies), and private debt. These components are not included in the computations after 1996. Until 1996 the value of stocks was self-reported, while afterwards it became third-party reported by banks and financial institutions (excluding non-listed stocks). The registration of the company value of self-employed has changed several times, but has stayed unchanged since 1997, where assets and liabilities of the firm were registered separately and included, respectively, in the assets and liabilities of the owner. Another definitional change occurs in 1983. Before 1983 all family wealth of a married couples was assigned to the husband, while the wealth of husbands and wives has been registered separately afterwards.

Ideally we would like to observe wealth, income, etc., of the individual in the middle age and observe the different generations at the same age because of the life-cycle variation in economic outcomes (Haider and Solon, 2006). If, for example, parents are around 25 years old when their children are born and we observe wealth of the child generation in 2011, then we would like to observe wealth of the parents and grandparents in 1986 and 1961, respectively. This goal has to be balanced against data availability (grandparents) and data quality (parents). Our main empirical analyses are based on parental wealth observed in 1997-1999, where the definition of the wealth measure is the same as that used for children in 2009-2011, and grandparental wealth measured in 1983-1985 where wealth of biological grandfathers and grandmothers are more accurately measured than the years before. We take three year averages of wealth of each individual to reduce the importance of transitory components, as often done in the literature following Solon (1992). The effects on our estimates of this procedure are rather small.

The largest change in the definition of wealth occurs around 1997 where the wealth tax was abolished. However, for 1995 and 1996 Statistics Denmark computed assets and liabilities of each individual using both the new definition of wealth (used for children and parents) and the old definition (used for grandparents). In Appendix A, we exploit this overlap to show that the new wealth measure is well approximated by the old way of measuring wealth, and we provide more details on the wealth data. Additional information on the Danish wealth and income-tax data may be found in Leth-Petersen (2010) and Chetty et al. (2011).

In the empirical analysis, we consider two types of samples: a child-parents sample (CP) and a child-parents-grandparents (CPG) sample. The CP sample focuses only on child-parents relationships without exploiting information of grandparents. In this sample, we consider all children of age 21-51 in 2011 (ensuring that they are born in 1960 or later), where both parents are alive in 2011, and where both parents are between 21 and 66 years old in 1999. This is very similar to Charles and Hurst (2003) with the exception that they limit the sample to children that are older than 25 years and require that both children and parents have positive wealth. The importance of these two differences in sample selection will be discussed. The child-parents-grandparents (CPG) sample is based on the CP sample, but with the additional requirement that at least one grandparent is alive in 1985. To avoid selection problems, we further require that parents are born in 1960 or later, corresponding to a maximum age of 39 in 1999, implying that the personal identifiers of all the grandparents are known.[6]

Table 1 provides summary statistics of the two samples. In the CP sample, we have 1.16 million child-parent pairs and the CPG sample consists of 97 thousand observations. In both samples, parents are significantly older than their children at the time where wealth levels are observed in the data, and grandparents are older than parents in the CPG sample. Since households normally accumulate wealth over the life cycle up to retirement, we should expect to observe the highest wealth for grandparents and higher wealth for parents than for children, which is also the pattern we see in Table 1. The large sample sizes allow us to account for the life-cycle effects by age-adjusting the wealth levels using age dummies and to estimate separate effects for different child cohort–parent age constellations in the empirical analyses. Moreover, following the existing literature measuring intergenerational effects, we will relate relative differences within a generation

[6]This reduces the sample considerably. Without the restriction, we obtain the same regression coefficients and a higher statistical precision but we prefer the restricted sample to avoid sample selection bias.

to the relative difference of another generation, which is not sensitive to scale effects.

Notice that wealth is negative for many individuals. This is in particular the case among the child generation where around half of the individuals have negative wealth compared to eight percent in the sample of C&H. One reason is that we do not restrict the sample to children above 25 as C&H, which significantly increases the wealth of children but also of parents, who will on average be older. Another reason is that Danish households have very high debt-to-income ratios (the liability-gross income ratio is around 200 percent for children and 150 percent for parents in the CP sample) compared to other countries, which has received international attention recently (IMF, 2012; European Commission, 2012). The difference to the US and other countries probably reflects that Denmark has a reasonably high universal public pension benefit level, substantial labor market pension savings by international standards, and an extensive social safety net that reduces the need for precautionary savings.

Labor earnings and gross income are on average higher for parents than for children, while earnings for grandparents are lower than for parents reflecting that some of the individuals have retired. The table also reports the years of education counting completed education. Parents are clearly more educated than grandparents but also somewhat more educated than their children in the CPG sample. However, this reflects that many of the children have not yet completed their education.

4 Empirical analysis

4.1 Elasticity of child wealth with respect to parental wealth

We first use the child-parents (CP) sample described above to estimate the (unconditional) elasticity of child wealth with respect to the average wealth of the (biological) parents.[7] This corresponds to estimating an ordinary least squares regression after using the natural log transformation on the wealth measures of children and parents. Column 1 of Table 2 reports the estimated elasticity without age adjustment, while column 2 reports the age-adjusted elasticity obtained by including age dummies of both children and parents in the regression. The finding of a child-parents age-adjusted elasticity equal to 0.268 implies that children born of parents with a wealth level that is 10 percent above

[7]Results are unchanged if we use the aggregate wealth of parents instead of the average wealth of parents, which is due to the fact that the log transformation and the IHS transformation of the data remove any scale effects.

the mean of the parent generation can expect to obtain a wealth level that is 2.7 percent above the mean of the child generation. C&H obtained an age-adjusted intergenerational wealth elasticity (IWE) of 0.365 for the United States using the PSID survey data. The lower estimate for Denmark is not surprising. Denmark has a very homogeneous population and a high degree of redistribution, and comparative studies have found that Denmark has the lowest intergenerational elasticity of earnings/income.[8]

When applying the log transformation, we are throwing away all child-parents pairs where either the child or the parents have zero or negative net wealth. Most of the empirical literature analyzing intergenerational relationships have looked at economic outcomes that do not attain negative values by definition, for example earnings. In this case, it is natural to apply the log transformation, which has appealing properties. This is, however, not the case when analyzing net wealth, which may well be negative, and where standard life cycle theory predicts negative values for young persons who have increasing earnings profiles. Another reason for observing negative wealth of households in our case, and also in C&H, is that we are unable to include pension wealth. In order to avoid the potential selection problem of using the log transformation, we will for the remaining part of the paper be using the inverse hyperbolic sine transformation (IHS), $w = \log(W + \sqrt{W^2+1})$, that behaves as $\pm\log(|W|)$ everywhere with the exception of in the neighborhood of zero.[9] Column 3 shows the IWE estimate after using the IHS transformation on the sample where wealth is restricted to be positive. The estimate is completely identical to the result based on the log transformation in column 2. Next, we consider the full sample with 1.16 million child-parents observations that include negative wealth of children and/or parents. This gives an IWE of 0.19, which is considerably lower than the estimate obtained from the restricted sample, showing that it may create severe selection bias to remove observations with negative wealth. Note finally from Table 2 that all regression coefficients are very precisely estimated because of the large sample size as illustrated by the tiny standard deviations of the estimates.[10]

[8] An overview of estimates of intergenerational earnings elasticities for different countries may be found in Björklund and Jännti (2009). They report an elasticity for Denmark of 0.12, which makes Denmark the country with the lowest correlation across generations.

[9] For details on why the IHS transformation can be used to estimate approximate elasticities, see Appendix B.

[10] Another concern may be that outliers or observations with zero or close to zero wealth may be very important for the estimates. We have run sensitivity analyses, which revealed no effects on the estimates of removing observations in the tale of the distribution and around zero wealth.

Life-cycle variation may be important when measuring intergenerational mobility in economic outcomes (Haider and Solon; 2006). Thus, there may be substantial variation in the IWE, depending on the age of the child and the parents, that are not removed by including age dummies. The large sample size allows to address this issue by running separate regressions for each age-cohort of the children and for each age-cohort of the parents measured at the year where the child was born, respectively. Figure 1a shows the IWE estimate as a function of the age of the child in 2011 starting from the early twenties and going up to an age of fifty years. The diagram is constructed by running separate regressions for each cohort of the children, including in each regression age dummies of the parents, and then plotting the estimates and the 95 percent confidence interval for each cohort. The confidence interval shows that the IWE is very precisely estimated for each age group and the graph displays a remarkably stable IWE. For all thirty age groups, estimates lie within a narrow interval from 0.16 to 0.22 without any systematic trend. Figure 1b is constructed in the same way but this time subsamples are created for each age level of the parents at the time when the child was born and child age dummies are included in each regression. The graph shows a weakly increasing correlation of wealth between children and parents when the age-difference between parents and children is increased, but the main conclusion is again that the IWE is remarkably stable and within the same interval as in Figure 1a.[11]

4.2 Decomposition of the intergenerational wealth elasticity

In Table 3, we add income controls. Column 1 is the baseline without income controls and is identical to column 4 in Table 2. We first include the income of the child. It is natural to expect that a high income level of the child is associated with a high level of wealth. This is also what we see from column 2 showing that the elasticity of child wealth with respect to child income equals 1.4, implying that a ten percent higher income level is associated with 14 percent higher wealth. Our main interest is how it influences the size of the IWE. If parents with high wealth invest more in the human capital formation of their children, as in the theory of Becker and Tomes (1979) then this would raise the income and wealth of their children. This would generate a positive correlation between

[11] We have also estimated the IWE for each child-parent age constellations. All estimates are significant with the exception of constellations where the age-difference between the child and the parents is very small or very large. The conclusion is again that estimates are roughly the same for all combinations.

child wealth and parental wealth working through child income, and if this channel was stronger than other sources of intergenerational correlation in wealth, the IWE estimate should fall when we introduce child income as a regressor. Table 2 shows that the IWE is unchanged after introducing child income, and does therefore not lend support to a mechanism working through child income.

Next, we also introduce parental income in the OLS regression. Column 2 shows that the elasticity of child wealth with respect to parental income is 2.2, which is therefore more important than the child's own income. This corresponds closely to the specifications we discussed in our theoretical section: the objective is to shut down income as a mechanism of intergenerational transmission of wealth, and it requires controlling for income of both parents and children.

It could be that parental wealth was not directly relevant for wealth of children, but own income was an important determinant of own wealth. If it was so, correlation of income would translate into unconditional correlation of wealth. Unconditional correlation between wealth of children and parents could also arise if the true relationship was from parental income to child wealth, because parental wealth works as a proxy for parental income.

In each of these cases, the IWE would fall when controlling for parental income, but we observe nearly the same estimate as in column 2. Altogether, the IWE falls only slightly when going from column 1 to column 3. Income explains less than 8 percent of the IWE, and this conclusion is completely unchanged if we use a more flexible specification with, for example, fourth degree polynomials in child income and parental income, respectively.

The last two columns, report the result of splitting the sample depending on parental age. This has a sizable impact on the coefficients on income but only a small impact on the IWE. Note that the coefficient on the child's own income becomes negative. This may reflect that the children are young in this group and that students, who currently have a low income (student benefits), come from a wealthy background and are wealthy themselves compared to young persons earning income in the labor market.

In Table 4, we include number of siblings dummies, years of schooling dummies, and financial composition dummies. The idea here is to test the independent importance of the corresponding mechanisms beyond their relationship to wealth.

Column 1 is the baseline with only age dummies included. In column 2, we include number of siblings dummies, so that we only exploit the variation within families of a given size to estimate the IWE. It is certainly possible that the number of siblings might affect one's own wealth separately from any relationship that it might have with parental wealth. For example, siblings may reduce incentive to save if they provide an implicit insurance or the number of siblings may influence traits that are determinants of wealth accumulation such as patience. Inclusion of sibling dummies has nearly no impact on the estimate suggesting that sibling considerations, if any, are adequately captured by parental wealth. Similarly, dummies for education length of both parents and children reduce the IWE only a little, and the effect is of the same magnitude as when introducing income controls (see column 3 of Table 3), suggesting that wealth is also a good proxy for this channel.

The largest reduction in the IWE arises when we introduce financial composition dummies for children and parents, i.e., dummies for homeownership, stock ownership, bonds ownership, and for being self-employed. This reduces the IWE by 25 percent, which is much more than what is explained by the income controls alone. Finally, when including all control variables the IWE becomes 0.134, which is a reduction of the IWE by less than 30 percent when starting from the baseline with only age controls. Thus, the explanatory variables explain only up to 30 percent of the IWE, which is considerably less than in C&H where income alone explains more than 50 percent of the IWE, and all variables together explain nearly 2/3 of the IWE.

4.3 Wealth across three generations

In this subsection, we use the child-parents-grandparents (CPG) sample, which enables us to analyse wealth across three generations. The children and parents are younger in the CPG subsample than in the full CP sample. For comparison, we therefore start by running a simple regression of child wealth with respect to parental wealth for this subsample. The result is reported in column 1 of Table 5 and gives an age-adjusted elasticity of 0.177, which is only a little smaller than the elasticity obtained when estimating the relationship on the full sample (column 4 of Table 2).

It is important to know whether the relationship between child wealth and parental wealth is stable over time/generations. There may be long run forces that reduce or increase the strength of the relationship, and recent research has documented substantial changes over the long run in top income shares, in the relative importance of capital and labor income at the top of the income distribution, and in the evolution of inheritance (Atkinson, Piketty and Saez, 2011; Piketty, 2011). In column 2, we report the relationship between parental wealth and grandparental wealth in the data set. The age-adjusted child-parents wealth elasticity from this estimation is 0.160, which is only slightly lower than the IWE when running the regression on the child-parents pairs in the sample, and the standard deviations on the estimates from these two regressions are small and completely identical. This indicates that the underlying intergenerational relationship is quite stable.

Next, we look separately at the correlation between children and grandparents. From a structural relationship only relating generation g to generation $g-1$ such as $w_g = \beta_0 + \beta_1 w_{g-1} + \varepsilon$, we would expect that a regression of w_g on w_{g-2} would give a coefficient equal to $(\beta_1)^2$. From the IWE estimates in columns 1 and 2, we should therefore expect a coefficient around 0.03 (i.e., $0.177^2 = 0.031$ and $0.160^2 = 0.026$). The estimated coefficient we obtain from regressing child wealth on grandparental wealth is 0.094 (column 3 of Table 5) and is therefore more than three times as large. One reason for this difference could be that the underlying stochastic processes relating wealth across generations have more memory than just one generation, for example because grandparents have a direct impact on their grandchildren (Solon, 2012). Another possible reason, which we pursue in the next subsection, is measurement error/omitted variable that creates a downward bias in the estimated coefficient.[12]

Column 4 of Table 5 reports the results from including both parental wealth and grandparental wealth in the estimation. When compared to the univariate relationships (columns 1 and 2), we see that the coefficient on parents fall a little, whereas the coefficient on grandparents falls by 1/3. If we consider a ten percent increase in the wealth of grandparents then this regression predicts a nearly one percent higher wealth of the grandchildren, which is again three times as high as the prediction we obtain from a standard estimation of IWE exploiting data from only two generations. In Table 6, we add the same controls as we did in Table 4 plus dummy variables for number of cousins

[12] If, for example, only 50 percent of the measured variation in wealth is governed by the true variation while the rest is noise then we would estimate a β_1- coefficient around 0.2, when the structural coefficient is 0.4. The true coefficient on grandparents would be $0.4\hat{}2 = 0,16$, but we would estimate $0.16 * 50\% = 0.08$, which is higher than $0.2\hat{}2 = 0.04$.

and for number of living grandparents in 2011.[13] Introducing the demography variables in isolation reduces the coefficients on parents and grandparents a little (going from column 1 to column 2). The effects on the coefficients are larger when including education and financial composition dummies (column 3), as was also the case when we looked at the relationship across only two generations in Table 4. The effect of grandparental wealth on child wealth is again three times as big as the predicted effect from a standard two-generation estimation of the IWE.

4.4 Using grandparental wealth as instrument

Attenuation bias caused my measurement problems has been a main issue in the intergenerational income mobility literature since the influential contribution by Solon (1992). Measurement problems may for example be caused by response errors in surveys and by transitory components in income implying that current income is a pure measure of the permanent income. The standard method to reduce the influence of the transitory component is to compute three or five year averages, as we have also done in our measurement of wealth. However, this method may only remove a small part of the attenuation bias if the transitory component has some persistence. For example, Mazumder (2005) provides simulations suggesting that it may require an average over 20 to 30 years in order to bring the attenuation factor down to 90 percent.

The theory in Section 2 illustrates cases where child wealth only depends on parental wealth and where ordinary least squares estimation of the relationship provides a downward biased estimate of the IWE, but where two-stage least squares estimation using wealth of grandparents as instrument for wealth of parents provides a consistent estimator. Table 7 reports results from such 2SLS estimations. The two first columns report 1st stage and 2nd stage results with only age dummies as additional regressors, while the following columns report results when other variables are added. The strong correlation of parental wealth with grandparental wealth and F-test values well above 10 in all cases indicate that we do not have a 'weak instrument' problem in the 1st stage regressions.[14]

[13] If grandparents transfer money to grandchildren then the effect of many cousins would be less money received on average per grandchild. By including dummy variables for number of cousins, we are only exploiting the variation within families of same size (measured by the number of cousins). We include dummy variables for number of living grandparents in 2011 because grandparents may die after we have observed their wealth in 1983-85, implying that parents and children may inherit their wealth.

[14] As a further test of the strength of the first stage, we have constructed a worst case scenario where grandparental wealth has no predictive power for parental wealth following the approach of Bound,

The 2nd stage results give estimates of the IWE in the range 0.6-0.7, which is more than three times as high as the ordinary least squares estimates described in Subsection 4.1. For robustness, we have redone the age-dependency graphs in Figure 1a and 1b, but this time displaying the 2SLS estimates. This is done in Figure 2a and 2b, which for comparison also include the OLS estimates of Figure 1a and 1b. The graphs based on the 2SLS estimates have the same shapes as the graphs based on the OLS estimates, but of course with a more wide confidence interval. More importantly, the IWE estimates are consistently around three times as high at each age level.

4.5 Dynasty fixed effects

In Section 2, we also consider the possibility of a dynasty/social class fixed effect leading to the relationship

$$w_g^d = b_0^d + b_1 w_{g-1}^d + \varepsilon_g, \quad g = 1, 2, \tag{8}$$

where w_g^d is the wealth level of generation g in dynasty d, b_0^d is the fixed effect of the dynasty and ε_g is an error term. A direct estimation of equation (8) will provide an inconsistent estimate of b_1 because the assumption of strict exogeneity of the regressors is violated by construction when using the lagged dependent variable as regressor. The Within estimator is also biased and so is an estimation of the first differences:

$$\Delta w_g^d = b_1 \Delta w_{g-1}^d + \Delta\varepsilon_g, \tag{9}$$

where $\Delta w_g^d \equiv w_g^d - w_{g-1}^d$. However, our assumption that ε_g is serially uncorrelated allows us to follow the empirical approach first suggested by Anderson and Hsiao (1982) and use grandparental wealth w_{g-2}^d as an instrument for Δw_{g-1}^d. This avoids the bias of the lagged dependent variable and provides a consistent estimate of b_1. Table 8 shows the results from this exercise where we have allowed each child-parents-grandparents to have a unique constant b_0^d.[15] The first column reports the OLS estimate of b_1 under the assumption of a common constant $b_0^d = b_0$ for d, while the second column reports the first-differences 2SLS estimate of b_1. The table shows that the 2SLS estimate of

Jaeger, and Baker (1995). We achieve this by randomly assigning the observations of grandparental wealth to the observations in data. The resulting pseudo instrument is by construction uncorrelated with the endogenous regressor, yet, it retains the marginal distribution of the actual instrumental variable. Using the pseudo-instrument gives an IV estimate close to the OLS estimate, and standard errors are much higher compared to the actual IV standard errors.

[15] We have chosen this completely flexible procedure rather than, for example, restricting siblings to have the same fixed effect because they have the same family background.

the coefficient is only around 1/3 of the OLS estimate. Thus, allowing for fixed effects instead of a common constant reduces the estimate substantially, strongly suggesting more persistence in wealth formation across generations and that the standard measurement of the IWE is underestimating the level of intergenerational fluidity.

4.6 Robustness

In the empirical analysis, we have imposed the condition that both parents of a child should be alive in 2011 and that some grandparents should be alive in 1985. We have therefore allowed for the possibility that grandparents die between 1985 and 2011 and leave wealth to their heirs, and we have allowed for variation in the number of grandparents alive in 1985. In order to study the importance of these sample selection criteria, we have redone the baseline regressions for the subsample where all four grandparents are alive in 2011. The results are shown in Table 9. The basic child-parents wealth elasticity equals 0.17 (0.19 before), the basic child-grandparents elasticity is 0.06 (0.09 before) and the 2SLS estimate of the child-parents wealth elasticity is 0.52 (0.61 before). The estimates are therefore somewhat smaller, which may reflect the presence of bequest in the large sample or just that individuals are younger in the subsample, but without changing the main conclusions.

5 Concluding remarks

Using administrative wealth records, we have estimated the child-parents IWE. The theory section illuminates that the IWE is generally not a deep parameter, but a complex function of the underlying mechanisms, by which the intergenerational transmission occurs.

However, we find a pronounced stability of the IWE across specifications, time, and generations, suggesting that parental wealth may serve as a sufficient statistic for the effect of parental characteristics on children's wealth that vary across different contexts that we consider. In particular, lack of substantial sensitivity to many controls (including income and education) suggests that the role of these specific channels in explaining wealth of children is captured by parental wealth.

Exploiting that Danish administrative wealth records allow us to use data for three generations, we find that there is role to be played by grandparental wealth in predicting their grandchildren's wealth, suggesting that the wealth generating process has longer memory than just one generation. As we discuss, an alternative interpretation of this may be that of measurement error in parental wealth, where grandparental wealth, under strong assumptions, may serve as an instrument for parental wealth, producing a much higher IWE compared to the basic estimate. On the other hand, grandparents may simply have a direct effect on grandchildren (invalidating the IV procedure). Either way, our empirical evidence illustrates the value added from having more generations of wealth data in characterizing the process that governs wealth across generations.

A Additional description of the wealth data

A.1 More details on the wealth data

The wealth data records are subject to a number of data breaks due to changing classifications of certain assets and liabilities, changes in reporting requirements, changes in tax treatment, etc. Table A.1 provides an overview of all data breaks in the underlying components of assets and liabilities. Although we do not construct the aggregate measures of assets and liabilities but instead rely on those compiled by Statistics Denmark, the table is still informative regarding the stability of the definition of wealth in the period. On the larger lines, subcomponents are generally third-party reported since 1997, whereas they relied in part on third-party reports and self-reporting prior to 1997.

The treatment of company values for self-employed has undergone a number of changes since the 1980s. In the period 1981 to 1985, firm assets such as buildings and operating fixture, equipment, machines, and cars are included. In 1981, buildings and operating fixture, equipment, machines, cars, etc. are registered at 80 pct. of the cash value or the balance sheet book value. It is calculated as 75 pct. of the cash value in 1982, 70 pct. of the cash value in 1983-1988, and 60 pct. of the cash value in 1988-1996. In the period 1986 to 1996, the equity of the firm is computed separately and included in the assets of a self-employed. In addition to the above assets, the computation of firm equity also includes financial assets of the firm, inventory, etc., and company debt is subtracted. From 1997 only firm assets are included in the assets of the owner while firm liabilities are included in the liabilities of the owner.

A.2 Analysing the impact of the 1997 change in the definition of wealth

The largest change in the definition of wealth occurs around 1997 where the wealth tax was abolished. However, for 1995 and 1996 Statistics Denmark computed assets and liabilities of each individual using both the new definition of wealth (used for children and parents) and the old definition (used for grandparents). In Table A.2 and Figures A.1 and A.2 we exploit this overlap to show that the new wealth measure is well approximated by the old way of measuring wealth.

In Table A.2 we run OLS regressions of the new definition of wealth measured in 1995 on the old definition of wealth measured also in 1995, and fixing the constant tern at zero. Columns 1 and 2 run regressions for grandparental and parental wealth, respectively, in the CPG sample, and column 3 runs the regression for parental wealth in the CP sample.[16] In all three cases, the slope coefficient is fairly close to unity, supporting our claim that the new definition approximates well the old definition.

Figure A.1 presents scatter plots of the new definition plotted against the old definition of wealth in 1995 for the same three cases as mentioned above. Axes are confined to wealth levels of no more than ± 1 million DKK (in 2010-prices) for readability. Again we find that observations tend to be concentrated around the 45 degree line through origo.

In Figure A.2 we show histograms of the deviations of the new definition of wealth in 1995 from the old defintion. The histograms exclude observations with no deviation and confine the primary axis to discrepancies of less than ± 1 million DKK (in 2010-prices) for readability. The distribution of discrepancies for all three cases are symmetric around zero and do not show any tendency for the new definition to systematically over- or underpredict old definition wealth.

B Estimating approximate elasticities with the inverse hyperbolic sine (IHS) transformation

The IHS transformation, defined as $w = IHS(W) = \log\left(W + \sqrt{W^2+1}\right)$, behaves as $\pm\log(|W|)$, except in a neighborhood around zero. Intuitively, we must then be able to estimate approximate elasticities in regressions of an IHS transformed variable on another

[16] Notice that the CP and CPG sample sizes are slightly lower than those reported in Table 1. This is due to observations with missing values in 1995, e.g., due to temporary emigration or, in the case of grandparents, death.

IHS transformed variable. To see this, first note that the derivative of the IHS function is $1/\sqrt{W^2+1}$. The slope parameter from regressing $y = IHS(Y)$ on $x = IHS(X)$ and a constant term is given by

$$\hat{\theta} = \frac{\partial E\left[y|x\right]}{\partial E\left[x\right]} = \frac{\partial Y}{\sqrt{Y^2+1}} \frac{\sqrt{X^2+1}}{\partial X} \approx \frac{\partial Y}{\partial X} \frac{|X|}{|Y|},$$

when X and Y are not too close to zero. In practice the approximation works well for wealth data, where the bulk of the mass of the wealth distribution is located away from zero.

References

[1] Anderson, T.W. and C. Hsiao (1982). "Formulation and Estimation of Dynamic Models Using Panel Data." *Journal of Econometrics* 18, 47–82.

[2] Arrondel, Luc (2009). "My Father Was Right: The Transmission of Values between Generations." PSE Working Paper 2009-12.

[3] Atkinson, Anthony B., Thomas Piketty and Emmanuel Saez (2011). "Top Incomes in the Long Run of History." *Journal of Economic Literature*, 49, 3–71.

[4] Becker, Gary and Nigel Tomes (1979). “An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility.” *Journal of Political Economy* 87, 1153 - 89.

[5] Becker, Gary and Nigel Tomes (1986). “Human Capital and the Rise and Fall of Families.” *Journal of Labor Economics* 4, 1-39.

[6] Björklund, Anders and Jäntti, Markus (2009). "Intergenerational Income Mobility and the Role of Family Background." in W. Salverda, B. Nolan and T. Smeeding, eds., *Oxford Handbook of Economic Inequality*, Oxford: Oxford University Press, 491-521.

[7] Black, Sandra and Paul J. Devereux (2011). "Recent Developments in Intergenerational Mobility." In: Ashenfelter, Orley, Card, David (Eds.), *Handbook of Labor Economics*. Vol. IVb. North-Holland, Amsterdam, 1487–1541.

[8] Bound, John, David A. Jaeger and Regina M. Baker (1995). "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogeneous Explanatory Variable is Weak." *Journal of the American Statistical Association* 90, 443-450.

[9] Charles, Kerwin K., Hurst, Erik (2003). "The Correlation of Wealth Across Generations." *The Journal of Political Economy* 111, 1155–1182.

[10] Chetty, Raj, John N. Friedman and Tore Olsen and Luigi Pistaferri (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.

[11] European Commission (2012). "Alert mechanism report." `http://ec.europa.eu/economy_finance/economic_governance/documents/alert_mechanism_report_2012_en.pdf`.

[12] Goldberger, Arthur S. (1989). "Economic and Mechanical Models of Intergenerational Transmission." *American Economic Review* 79, 504-513.

[13] Haider, Steven and Gary Solon (2006). "Life-Cycle Variation in the Association between Current and Lifetime Earnings." *American Economic Review* 96, 1308–1320.

[14] International Monetary Fund (2012). "World Economic Outlook (WEO): Growth Resuming, Dangers Remain." April 2012 (Washington).

[15] Kleven, Henrik Jacobsen, Martin Knudsen, Claus Thustrup Kreiner, Søren Pedersen and Emmanuel Saez (2011). "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark." *Econometrica* 79, 651–692.

[16] Kreiner, Claus Thustrup, David Dreyer Lassen and Søren Leth-Petersen (2012). "Measuring the Accuracy of Survey Responses using Administrative Register Data: Evidence from Denmark." Manuscript prepared for "*Improving the Measurement of Household Consumption Expenditures*," eds. Christopher Carroll, Thomas Crossley and John Sabelhaus. Forthcoming NBER Book Series Studies in Income and Wealth.

[17] Kolodziejczyk, Christophe (2011). "The Intergenerational Correlation of Wealth and Portfolio Allocation." Unpublished working paper.

[18] Leth-Petersen, Søren (2010). "Intertemporal Consumption and Credit Constraints: Does Consumption Respond to An Exogenous Shock to Credit?" *American Economic Review*, 100(3), pp. 1080-1103.

[19] Mazumder, Bhashkar (2005). "Fortunate sons: new estimates of intergenerational mobility in the US using social security earnings data." *Review of Economics and Statistics* 87, 235-255.

[20] Piketty, Thomas (2000). "Theories of Persistent Inequality and Intergenerational Mobility." In: Atkinson, A.B. and F. Bourguignon (Eds.), *Handbook of Income Distribution.* Vol I. North-Holland, Amsterdam, 430–476.

[21] Piketty, Thomas (2011). "On the Long-Run Evolution of Inheritance: France 1820–2050." *Quarterly Journal of Economics* 126, 1071-1131.

[22] Solon, Gary (1992). "Intergenerational income mobility in the United States." *American Economic Review* 82, 393–408.

[23] Solon, Gary (1999). "Intergenerational mobility in the labor market." In: Ashenfelter, Orley, Card, David (Eds.), *Handbook of Labor Economics.* Vol III. North-Holland, Amsterdam, 1761–1800.

[24] Solon, Gary (2012). "Theoretical Models of Inequality Transmission across Multiple Generations." Working paper. Michigan State University.

[25] Stiglitz, Joseph E. (1969). "Distribution of Income and Wealth Among Individuals." *Econometrica* 37, 382-397.

Table 1: Summary Statistics

	Child-Parent (CP) Sample		Child-Parent-Grandparent (CPG) Sample		
	Children 2009-2011	Parents 1997-1999	Children 2009-2011	Parents 1997-1999	Grandparents 1983-1985
Age	33.9	48.5	23.4	35.0	47.1
	(8.2)	(7.6)	(2.3)	(2.2)	(5.1)
Years of education[a]	12.9	12.3	11.0	12.6	8.9
	(2.4)	(2.6)	(1.7)	(1.9)	(2.3)
Labor income	256,025	288,873	129,398	282,080	247,334
	(207,597)	(174,194)	(103,486)	(135,885)	(133,734)
Gross income	326,867	396,203	172,004	358,322	371,134
	(233,078)	(345,531)	(92,620)	(127,748)	(163,437)
Share owning stocks	0.22	0.26	0.15	0.14	0.07
	(0.41)	(0.36)	(0.36)	(0.28)	(0.14)
Share owning property[b]	0.47	0.59	0.11	0.59	0.41
	(0.50)	(0.34)	(0.32)	(0.41)	(0.23)
Share owning bonds[c]	0.06	0.06	0.03	0.01	0.11
	(0.23)	(0.19)	(0.16)	(0.08)	(0.20)
Share self-employed[d]	0.04	0.08	0.01	0.06	0.11
	(0.19)	(0.20)	(0.10)	(0.17)	(0.16)
Value of assets	655,952	923,242	95,113	584,975	828,626
	(1,988,246)	(2,413,591)	(367,708)	(883,319)	(1,094,917)
Value of liabilities	636,575	596,021	129,221	605,768	550,996
	(1,702,573)	(1,207,167)	(343,321)	(714,934)	(822,594)
Net wealth	19,377	327,220	-34,108	-20,793	277,894
	(1,117,250)	(1,928,688)	(215,500)	(536,283)	(701,004)
Percentiles of wealth					
20th	-232,694	-89,456	-90,626	-185,697	-10,472
40th	-68,913	47,364	-24,166	-87,330	111,541
60th	8,240	255,708	3,969	-14,398	273,039
80th	147,696	585,444	28,575	123,348	499,783
Observations	1,155,564	1,155,564	97,438	97,438	97,438

Notes: The table reports mean values and standard deviations (in parentheses) of the variables. Age, education and ownership variables are as of 2011 for children, 1998 for parents and 1983 for grandparents.

Child-Parent (CP) sample: Children are aged 21-51 in 2011, both parents are alive in 2011 and aged 21-66 in 1999, and children are neither immigrants nor descendants of immigrants.

Child-Parents-Grandparents (CPG) sample: Children are aged 19-51 in 2011, both parents are alive in 2011 and aged 21-39 in 1999, children are neither immigrants nor descendants of immigrants, and have at least one grandparent alive in 1983. Parent variables are averages of biological parents. Grandparent variables are averages of biological grandparents. All monetary variables are measured in DKK and deflated with 2010 prices.

a) Measures years of completed education. The variable is based on 2010 data for children.

b) Property ownership dummy for grandparents is based on 1987 data.

c) Bond ownership dummy for grandparents is based on 1995 data.

d) Self-employed dummy for children is based on 2010 data.

Table 2: Child-Parent Wealth Elasticity

	(1) Child wealth 2009-2011 log	(2) Child wealth 2009-2011 log	(3) Child wealth 2009-2011 IHS	(4) Child wealth 2009-2011 IHS
Parental per cap. wealth (1997-1999, log)	0.379	0.268		
	(0.002)	(0.002)		
Parental per cap. wealth (1997-1999, IHS)			0.268	0.190
			(0.002)	(0.001)
Child age dummies		X	X	X
Parental age (avg., rounded) dummies		X	X	X
Observations	385,338	385,338	385,338	1,155,564
R-squared	0.086	0.267	0.268	0.102
Adj. R-squared	0.086	0.267	0.267	0.102

Notes: All regressions are on the CP sample described in Table 1. Difference in number of observations from (1)-(2) to (4) is that the log transform in (1)-(2) excludes observations of zero or negative wealth. The regression in (3) is run on the same sample as in (1)-(2). Robust standard errors are reported in the parentheses.

Table 3: Child-Parent Wealth Elasticity and the Importance of Own and Parental Income.

	(1) Baseline	(2)	(3)	(4) Par. age <= 50	(5) Par. age > 50
Parental per cap. wealth 1997-1999	0.190	0.186	0.175	0.165	0.185
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Child income 2009-2011		1.354	1.134	-0.299	2.772
		(0.022)	(0.022)	(0.027)	(0.035)
Parental per cap. income 1997-1999			2.215	3.057	1.276
			(0.029)	(0.041)	(0.042)
Child age dummies	X	X	X	X	X
Parental age (avg., rounded) dummies	X	X	X	X	X
Observations	1,155,564	1,155,564	1,155,564	611,918	543,646
R-squared	0.102	0.105	0.110	0.147	0.067
Adj. R-squared	0.102	0.105	0.110	0.147	0.067

Notes: Dependent variable is the average 2009-2011 wealth of children. The IHS transformation is used on wealth and income variables. All regressions are on the CP sample described in Table 1. Robust standard errors are reported in the parentheses.
Columns (4) and (5) split the sample by median parental age (50 in 1999).

Table 4: Child-Parent Wealth Elasticity and the Importance of Controls.

	(1)	(2)	(3)	(4)	(5)	(6)
Parental per cap. wealth 1997-1999	0.190	0.183	0.173	0.141	0.133	0.134
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Child income 2009-2011						-0.407
						(0.023)
Parental per cap. income 1997-1999						0.671
						(0.033)
No. of siblings dummies[a]		X			X	X
Years of schooling dummies[b]			X		X	X
Financial composition dummies[c]				X	X	X
Child age and parental age dummies	X	X	X	X	X	X
Observations	1,155,564	1,155,564	1,155,564	1,155,564	1,155,564	1,155,564
R-squared	0.102	0.105	0.111	0.155	0.161	0.162
Adj. R-squared	0.102	0.105	0.111	0.155	0.161	0.162

Notes: Dependent variable is the average 2009-2011 wealth of children. The IHS transformation is used on wealth and income variables. All regressions are on the CP sample described in Table 1. Robust standard errors are reported in the parentheses.
a) No. of siblings is calculated as the average of mother's no. of children and father's no. of children (rounded).
b) Years of schooling dummies for both child and parent are included. Years of schooling is recorded as of 2010.
c) Financial composition consists of dummies for both child and parents for homeownership, stock ownership, bonds ownership, and for being self-employed. All dummies are included for both child and parents. Parents are noted as owning stocks, say, if at least one parent owns stocks. Self-employment status for children is taken in 2010.

Table 5: Child-Parent and Child-Grandparent Wealth Elasticities

	(1) Child wealth 2009-2011	(2) Parental per cap. wealth 1997-1999	(3) Child wealth 2009-2011	(4) Child wealth 2009-2011
Parental per cap. wealth 1997-1999	0.177			0.168
	(0.003)			(0.003)
Grandparental per cap. wealth 1983-1985		0.160	0.094	0.062
		(0.003)	(0.003)	(0.003)
Child age dummies	X		X	X
Parental age (avg., rounded) dummies	X	X		X
Grandparent age (avg., rounded) dummies		X	X	X
Observations	97,438	97,438	97,438	97,438
R-squared	0.194	0.055	0.161	0.200
Adj. R-squared	0.194	0.055	0.160	0.199

Notes: The IHS transformation is used on wealth variables. All regressions are on the CPG sample described in Table 1. Robust standard errors are reported in the parentheses.

Table 6: Child-Parent and Child-Grandparent Wealth Elasticities, Including Various Control Variables.

	(1) Child wealth 2009-2011	(2) Child wealth 2009-2011	(3) Child wealth 2009-2011
Parental per cap. wealth 1997-1999	0.168	0.161	0.126
	(0.003)	(0.003)	(0.003)
Grandparental per cap. wealth 1983-1985	0.062	0.054	0.038
	(0.003)	(0.003)	(0.003)
No. of siblings dummies[a]		X	X
Number of cousins dummies[b]		X	X
Number of living grandparents in 2011 dummies[c]		X	X
Years of schooling dummies[d]			X
Financial composition dummies[e]			X
Full set of age dummies	X	X	X
Observations	97,438	97,438	97,438
R-squared	0.200	0.206	0.256
Adj. R-squared	0.199	0.205	0.255

Notes: The IHS transformation is used on wealth and income variables. All regressions are on the CPG sample described in Table 1. Robust standard errors are reported in the parentheses.
a) No. of siblings is calculated as the average of mother's no. of children and father's no. of children (rounded).
b) No. of cousins is calculated as the average of the (at most) four grandparents' no. of grandchildren (rounded).
c) Al combinations of four dummies (one for each grandparent) denoting whether a grandparent is alive in 2011.
d) Years of schooling dummies for both child and parent are included. Years of schooling is recorded as of 2010.
e) Financial composition consists of dummies for both child and parents for homeownership, stock ownership, bonds ownership, and for being self-employed. All dummies are included for both child and parents. Parents are noted as owning stocks, say, if at least one parent owns stocks. Self-employment status for children is taken in 2010.

Table 7: Child-Parent Wealth Elasticities, 2SLS Estimation Using Grandparental Wealth As Instrument.

	(1) 1st stage IV	(2) 2nd stage IV	(3) 1st stage IV	(4) 2nd stage IV	(5) 1st stage IV	(6) 2nd stage IV
Grandparental per cap. wealth 1983-1985	0.155		0.151		0.087	
	(0.003)		(0.003)		(0.003)	
Parental per cap. wealth 1997-1999		0.608		0.603		0.631
		(0.022)		(0.023)		(0.040)
Number of siblings dummies			X	X	X	X
Number of cousins dummies			X	X	X	X
Extra control variables*					X	X
Observations	97,438	97,438	97,438	97,438	97,438	97,438
R-squared	0.067	.	0.074	.	0.182	0.016
Adj. R-squared	0.067	.	0.073	.	0.181	0.015
F-test (1st stage)	2,097.88		1,976.82		700.42	
F-test p.-val. (1st stage)	0.0000		0.0000		0.0000	
Partial R-sq (1st stage)	0.0186		0.0176		0.0064	
R-sq (1st stage)	0.0670		0.0739		0.1819	
Adj. R-sq (1st stage)	0.0667		0.0732		0.1809	

Notes: The IHS transformation is used on wealth and income variables. All regressions are on the CPG sample described in Table 1. Robust standard errors are reported in the parentheses. All regressions include both child and parental age dummies.
*The extra control variables used are identical to the ones in regression (6) of Table 4.

Table 8: Wealth Elasticities With Child Fixed Effects.

	(1) OLS	(2) FD 2SLS
Parental per cap. Wealth 1997-1999	0.220	0.080
	(0.003)	(0.005)
Observations	97,438	97,438
R-squared	0.055	.

Notes: Dependent variable is the average 2009-2011 wealth of children. The IHS transformation is used on wealth variables. All regressions are on the CPG sample described in Table 1. Robust standard errors are reported in the parentheses. Estimation details: OLS estimation of child wealth regressed on parental wealth and a constant. FD 2SLS is the difference between child and parents regressed on the difference between parents and grandparents, where the difference between parents and grandparents is instrumented using grandparental (i.e., double lagged) per capita wealth.

Table 9: Child-Parent Wealth Elasticities, Robustness Check Using Families With 4 Living Grandparents.

	(1) OLS	(2) OLS	(3) 1st stage IV	(4) 2nd stage IV
Parental per cap. wealth 1997-1999	0.165			0.521
	(0.006)			(0.062)
Grandparental per cap. wealth 1983-1985		0.059	0.137	
		(0.008)	(0.009)	
Child age dummies	X	X	X	X
Parental age (avg., rounded) dummies	X		X	X
Grandparent age (avg., rounded) dummies		X		
Observations	24,383	15,091	15,091	15,091
R-squared	0.193	0.166	0.069	0.042
Adj. R-squared	0.192	0.163	0.068	0.040
F-test (1st stage)			246.61	
F-test p.-val. (1st stage)			0.0000	
Partial R-sq (1st stage)			0.0147	
R-sq (1st stage)			0.0692	
Adj. R-sq (1st stage)			0.0675	

Notes: The IHS transformation is used on wealth variables. The results of the regression in column 1 correspond to column 4 of table 2 but now using only families with 4 living grandparents in 2011. Likewise the results in column 2 correspond to column 3 of table 5 and results in column 3 and 4 correspond to column 1 and 2 of table 7. Robust standard errors are reported in parentheses.

Table A.1: Wealth Data Documentation

	Year																																
	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10	11	Variables
Assets																																	(*=ends with akt)
Bank deposits					*a*								*b*	*c*				*d*															INDESTPI/ BANK*
Bonds					*a*								*e*					*f*				*g*											KURSOBL/ INDESTPI/ OBL*
Deeds					*a*								*h*	*e*				*d*															KURSPANT/ INDESTPI/ PANT*
Stocks					*a*										*i*			*d*				*g*											KURS*
Value of property				*j*	*a*													*k*															KOEJD/ KOEJD_NY05
Liabilities																																	(*=ends with GALD, GAELD or GÆLD)
Bank debt					*a,l*		*m*	*n*			*n*	*n*		*n*	*o*			*d*															GÆLDIO/ BANK*/ PRI*/ BANK*
Mortgage debt					*a,o*													*d*															PRI*/ OBL*
Deed debt					*a,o*			*p*					*q*	*r*				*d*															PRI*/ PANT*/ BANK*/ PANT*

Color coding:
- Combined 3rd party reporting and self reporting
- 3rd party reports
- Variable overlap
- Variable shift

Data breaks

a During 1984-1986, not incl. people with special income; Income from stocks and bonds, self-employed, or income from other countries. From 1987 and onwards everyone is included.
b Including market value of bonds.
c Including deeds in deposits.
d Wealth tax is abolished. Only 3rd party info. from financial institutions and banks is available.
e Included in the aggregate variable indestpi. No seperate registration.
f Wealth tax is abolished. Only 3rd party info. from financial institutions and banks is available. Only incl. bonds in deposits.
g From 2001 the part of a mutual fond placed in bonds is moved from KURSAKT to OBLAKT.
h Only incl. deeds in deposits.
i Excl. unlisted stocks from 1994-1996.
j Value of property is for 500,000 people only registered in the aggregate asset variable, QAKTIVF.
k KOEJD only includes residential property, KOEJD_NY05 also includes business property and undeveloped land.
l Included in the aggregate variable GÆLDIO, which includes all other debt (including debt to foreign countries) than mortgage debt. No separate registration.
m Independent registration of debt to financial institutions. Deeds issued by financial institutions is included.
n During 1987-1993 debt in own business is not included in BANKGAELD. From 1990 and onwards including debt in The Mortgage Bank, pension funds, insurance and financial companies, credit and debit card debt. From 1991 including student debt in financial institutions. From 1993 including debt in deposited deeds (PANTGALD). The variable does not exist in 1994, bank debt is only registrered as part of the aggregate variable PRIGÆLD.
o No separate registration. Variable included in the aggregate variable PRIGÆLD.
p During 1987-1993 deed debt in own business is not included in PANTGALD.
q From 1992 excluding debt in deeds, which are not deposited.
r Debt in deposited deeds is included in the variable BANKGAELD and later PRIGÆLD. No separate registration.

Table A.2: Wealth Data Break 1995, New Definition Regressed On Old Definition.

	CPG sample		CP sample
	(1)	(2)	(3)
	Grandparental wealth	Parental wealth	Parental wealth
Grandparental wealth	0.923		
	(0.001)		
Parental wealth		0.889	0.875
		(0.001)	(0.000)
Observations	97,328	97,388	1,155,090
R-squared	0.855	0.794	0.771
Adj. R-squared	0.855	0.794	0.771

Notes: Dependent variable is the new definition of wealth. Regressions in column 1-2 are on the CPG sample described in Table 1. The regression in column 3 is on the CP sample described in Table 1. The difference in sample size compared to Table 1 is due to missing values of 1995 wealth. The IHS transformation is used on wealth variables.
The constant term is fixed at zero. Robust standard errors are reported in the parentheses.

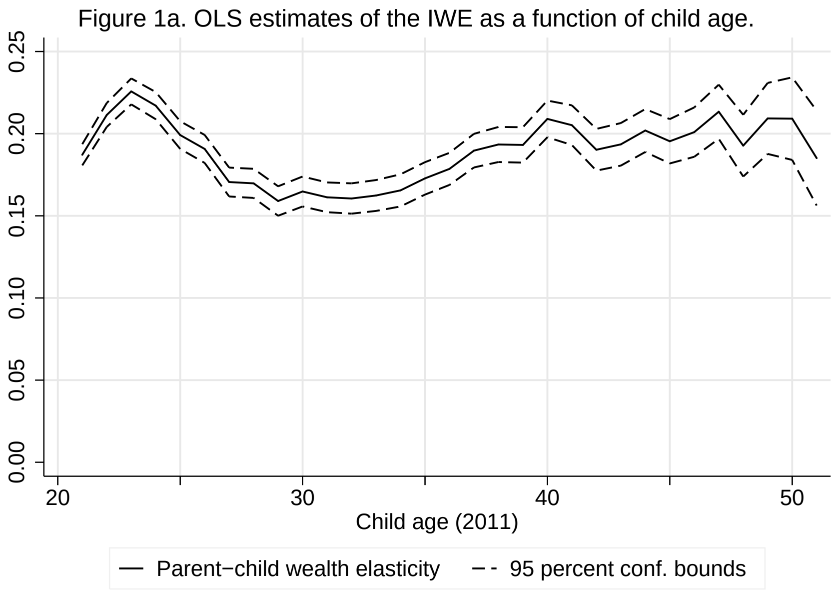

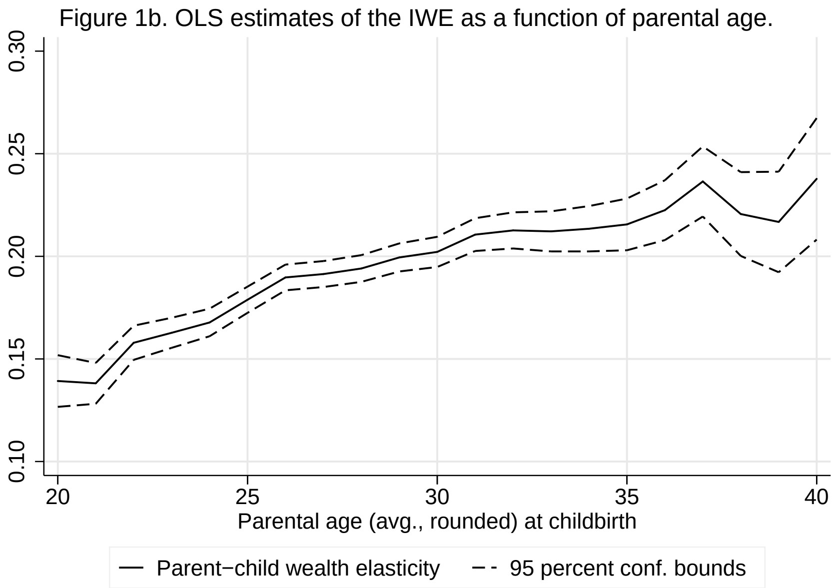

Figure 1: OLS Estimates of the IWE as a Function of (a) Child Age (b) Parental Age at Childbirth.

Notes: In Panel (a), estimates of the parent-child elasticity of wealth stem from separate regressions of child wealth (2009-2011, IHS) on parental per cap. wealth (1997-1999, IHS) and parental age (avg., rounded) dummies by child age in 2011. In Panel (b), estimates of the parent-child elasticity of wealth stem from separate regressions of child wealth (2009-2011, IHS) on parental per cap. wealth (1997-1999, IHS) and child age dummies by age of parents (avg., rounded) at childbirth. The regressions are run on the CP sample described in Table 1. 95 percent confidence bounds are calculated using robust standard errors.

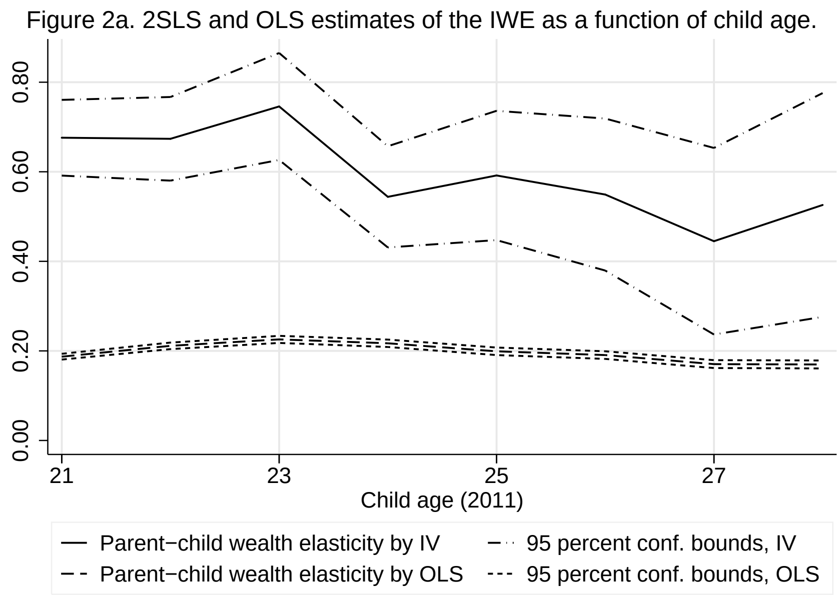

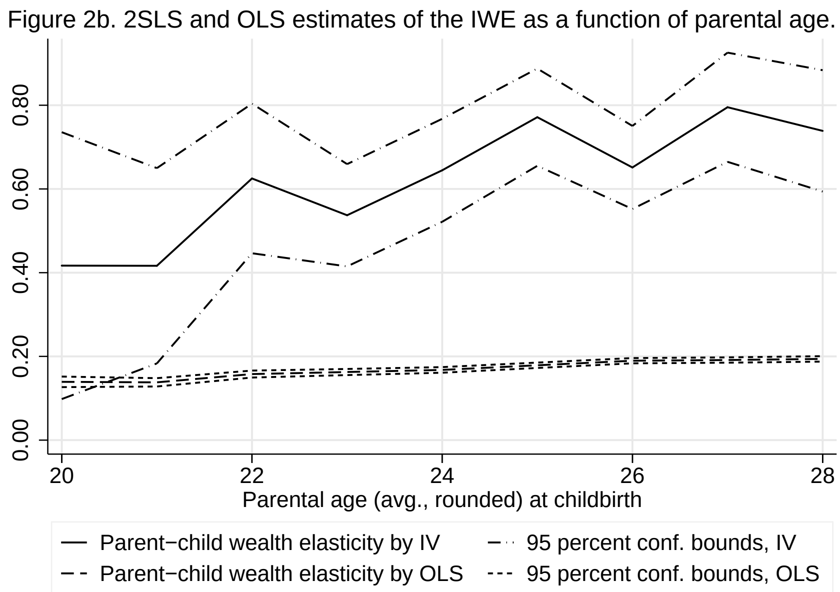

Figure 2: 2SLS and OLS Estimates of the IWE as a Function of (a) Child Age (b) Parental Age at Childbirth.

Notes: In Panel (a), estimates of the parent-child elasticity of wealth stem from separate regressions of child wealth (2009-2011, IHS) on parental per cap. wealth (1997-1999, IHS) and parental age (avg., rounded) dummies by child age in 2011 using 2SLS and OLS. In Panel (b), estimates of the parent-child elasticity of wealth stem from separate regressions of child wealth (2009-2011, IHS) on parental per cap. wealth (1997-1999, IHS) and child age dummies by age of parents (avg., rounded) at childbirth using 2SLS and OLS. In both panels, as instruments for parental wealth in 2SLS regressions, we use grandparental wealth (1983-1985, IHS) and a second degree polynomial in grandparental age (avg., rounded). OLS regressions are run on the CP sample (described in Table 1), and the 2SLS regressions are run on the CPG sample. 2SLS estimates are shown for child ages for which the first stage regression is strong, meaning that the F-test for joint exclusion of the instruments has an F-value above 10 and a p-value below 0.05. OLS estimates are shown for the same ages. 95 percent confidence bounds are calculated using robust standard errors.

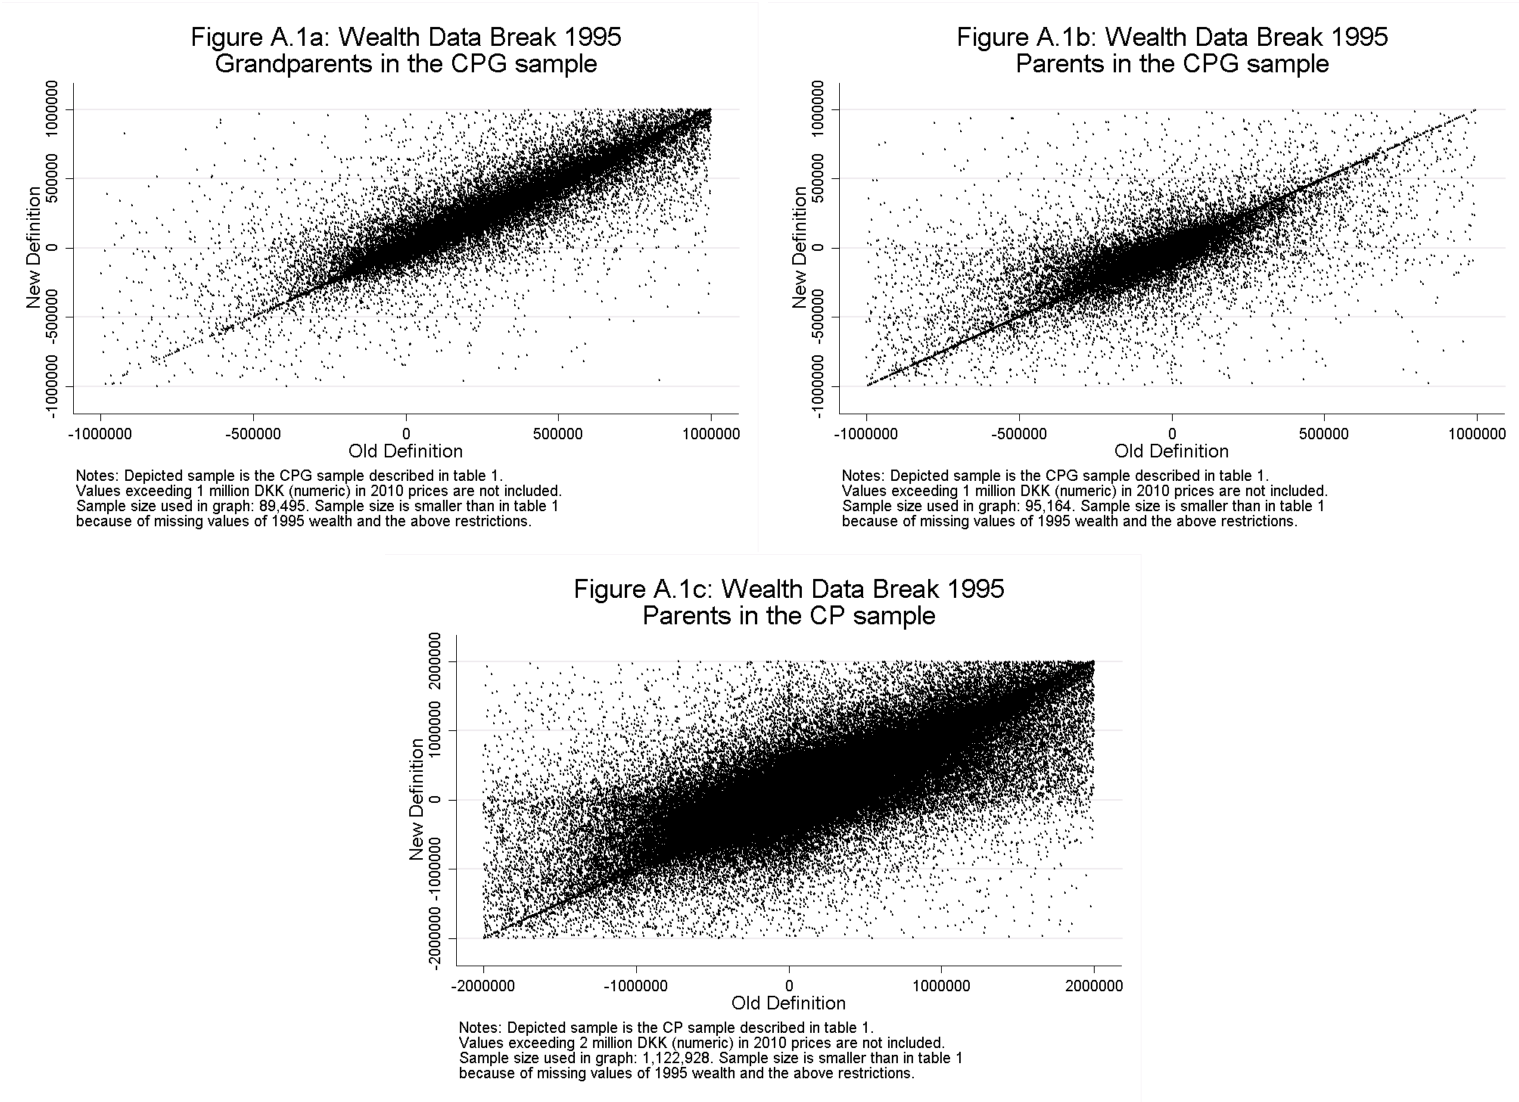

Figure A.1: Analysis of the Wealth Data Break in 1995.

Notes: Scatter plots of individual wealth as recorded by the new definition plotted against wealth as recorded by the old definition. Panel (a) and (b) show the plot for grandparents and parents, respectively, in the CPG sample. Panel (c) shows the plot for parents in the CP sample. Sample sizes differ compared to those in Table 1 due to missing observations in 1995, e.g. grandparents dying before 1995, and as we only show observations not exceeding 1 million DKK (in 2010-prices).

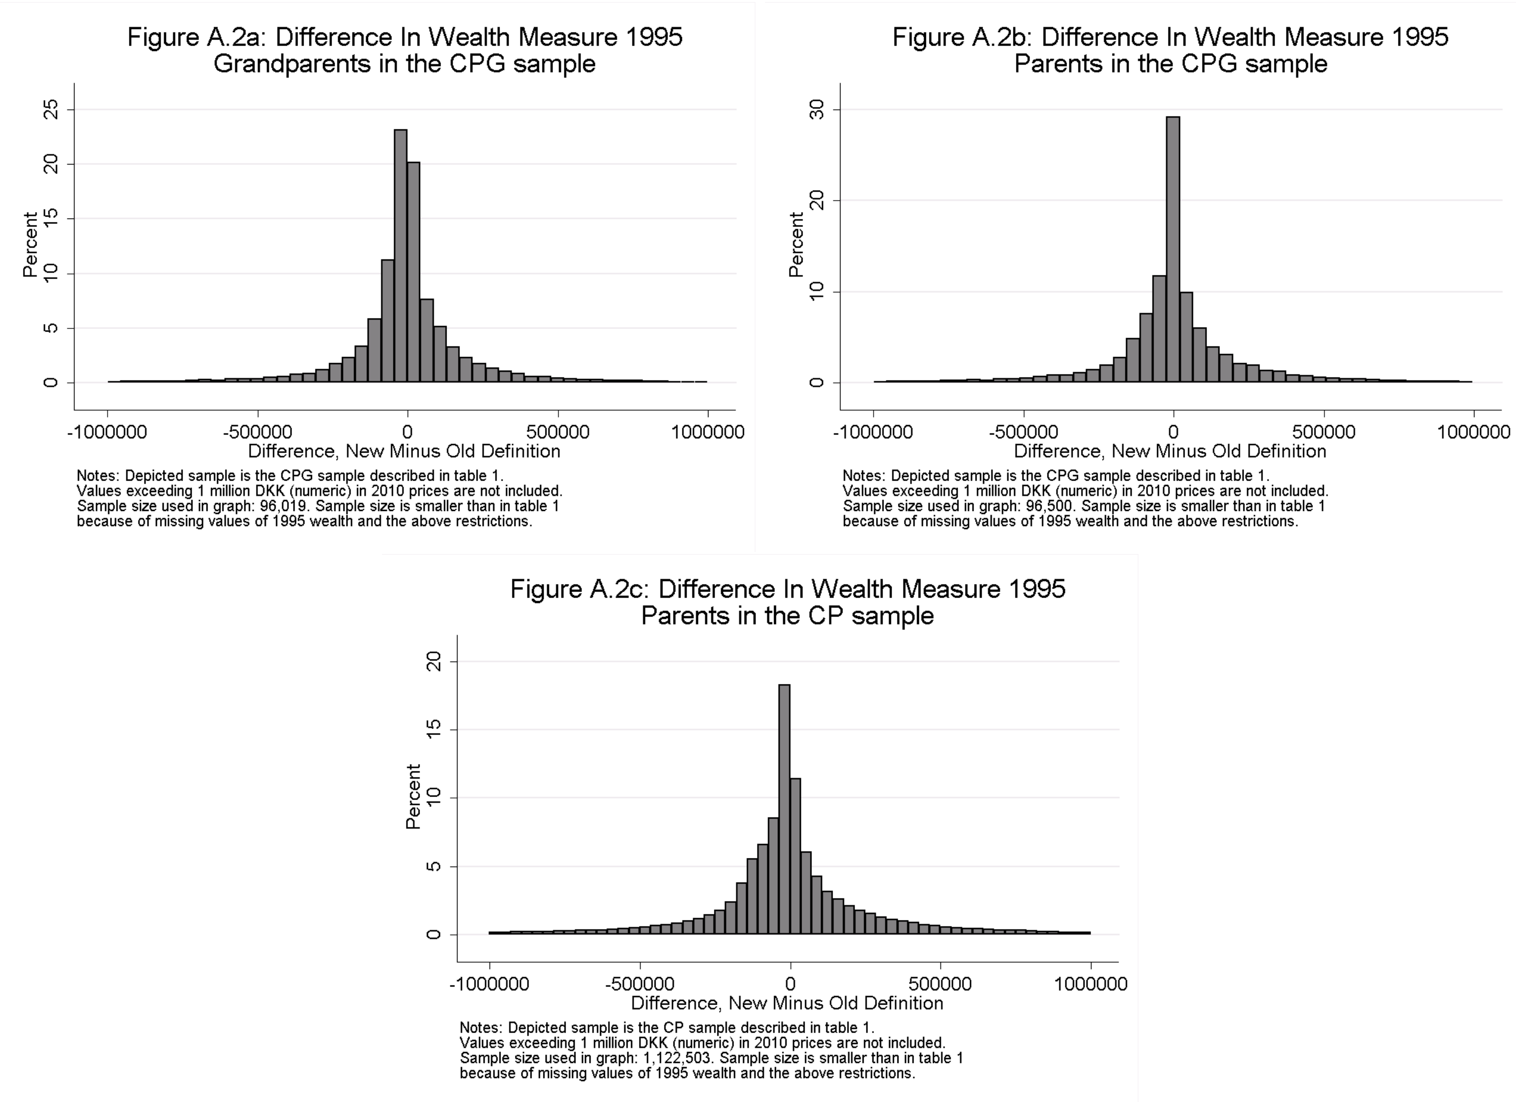

Figure A.2: Analysis of the Wealth Data Break in 1995.

Notes: Histograms of the difference in wealth as measured by the new and old definitions. Panel (a) and (b) show the histograms for grandparents and parents, respectively, in the CPG sample. Panel (c) shows the plot for parents in the CP sample. Sample sizes differ compared to those in Table 1 due to missing observations in 1995, e.g. grandparents dying before 1995, as we only show observations with discrepancies not exceeding 1 million DKK (in 2010-prices), and as we likewise exclude observations with no discrepancy.