

Verdelin, Nicolaj

Doctoral Thesis

Topics in public economics: Taxation, provision, and economic efficiency

PhD Series, No. 133

Provided in Cooperation with:

University of Copenhagen, Department of Economics

Suggested Citation: Verdelin, Nicolaj (2009) : Topics in public economics: Taxation, provision, and economic efficiency, PhD Series, No. 133, ISBN 978-87-91342-53-0, University of Copenhagen, Department of Economics, Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/233397>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Ph.D. Thesis

Topics in Public Economics:
Taxation, Provision, and Economic Efficiency

Nicolaj Verdelin

August 2008

Department of Economics

University of Copenhagen

Contents

	Acknowledgements	1
	Summary	3
	Summary in Danish	9
1	Intertemporal Tax Wedges and Marginal Deadweight Loss	15
2	Public Provision of Private Goods: Optimal Provision Schemes	51
3	Optimal Provision of Public Goods: A Synthesis	88
4	Taxation, Regulation, and Cue-Triggered Consumption	118
5	An Evaluation of the Tax-Transfer Treatment of Married Couples in European Countries	156

Acknowledgements

I would like to thank my supervisors Henrik Jacobsen Kleven and Claus Thustrup Kreiner for their advice and guidance throughout my time at the Department. I have always felt that Henrik and Claus have taken a genuine interest in my progression as a researcher and that they have persistently found the time to listen to my ideas and concerns. I have also had the pleasure of co-authoring with both Claus (twice) and Henrik and their enthusiasm and sharp minds have been a source of great inspiration. I am also grateful to them for initially encouraging me to pursue a Ph.D.

I also wish to thank my other co-authors, Herwig Immervoll and Jes Winther Hansen. Jes and I complement each other quite well and I have enjoyed working with him. Another source of joy has been my fellow Ph.D. students. I have appreciated the occasional Friday beer and the grim sense of humor and self-ironi when things were stalling.

I am also grateful to MIT for their hospitality when I visited there in 2005-2006. The atmosphere at the 'Institute' was a great inspiration and one of the most fulfilling learning experiences of my life. I am thankful to the people at MIT who made my stay such a pleasure and to my sponsors who made it possible.

I was fortunate to work as a research assistant at EPRU before I began writing my thesis. I learned a lot from this experience and there is no doubt that my time there was instrumental for my decision to pursue a Ph.D. I would like to thank Claus, Henrik, David Dreyer Lassen, Peter Birch Sørensen, and Hans Jørgen Whitta-Jacobsen for the opportunities they gave me as a research assistant. I also extend my thanks to the other

Acknowledgements

members of EPRU and everyone else at the Department.

Finally, I would like thank my family, especially my parents, and friends. They have been a constant source of support and wonderful distraction, providing a safe haven free from economics.

August 2008

Nicolaj Verdelin

Summary

This thesis consists of five self-contained chapters, which can be read independently. All chapters deal with the role of the government in the economy. Chapter 1 re-considers the cost of progressivity of the tax-transfer schedule when intertemporal responses are taken into account. Chapter 2 deals with the public provision of private goods and asks whether universal or means-tested provision is the better alternative. Chapter 3 is concerned with the optimal supply of public goods when the government relies on possibly distortionary income taxation. Chapter 4 uses insights from behavioral economics to study the spillovers between taxation and regulation of cue-triggered consumption goods. The final chapter, Chapter 5, evaluates the tax-transfer treatment of married couples in Europe based on a large microsimulation model. Apart from Chapter 5, which is entirely applied, all chapters contain theoretical contributions. Chapter 1 also applies the theoretical insights to U.S. data.

Chapter 1, *Intertemporal Tax Wedges and Marginal Deadweight Loss*, is written jointly with Jes Winther Hansen. In this chapter, we re-examine the efficiency of tax reforms in an explicitly dynamic setting. When income is taxed annually under a progressive tax schedule individual marginal tax rates vary from year to year. The resulting intertemporal tax wedges give taxpayers an incentive to shift income to years when they face low marginal tax rates.

Traditional assessments of the deadweight loss of income taxation have used a static framework that ignores intertemporal substitution and may have produced biased

estimates. Using a fully dynamic model, we derive an expression for the marginal dead-weight loss of income taxation entirely in terms of empirically observable elasticities and quantities. The expression can be separated into a distortion arising from the level of marginal tax rates and a distortion due to intertemporal tax wedges. We show that ignoring intertemporal substitution will lead to an underestimate of the marginal dead-weight loss of a tax reform that increases intertemporal tax wedges and an overestimate when intertemporal tax wedges decrease. In some cases, a marginal tax increase can even improve overall efficiency if the reduction in intertemporal tax wedges is sufficiently strong.

In order to quantify the effects of ignoring intertemporal substitution in taxable income we analyze three stylized tax reforms on a U.S. panel data set with about 34,000 nationally representative tax units. Our simulations are performed for a large range of behavioral elasticities as a reflection of the uncertainty that exists in the literature on the size of the intertemporal elasticity of substitution. Overall, we find substantial biases from ignoring intertemporal responses even for small elasticity values. These results suggest that tax progressivity may be much more costly than previously thought.

Chapter 2, *Public Provision of Private Goods: Optimal Provision Schemes*, is concerned with the design of the modern welfare state. Specifically, the chapter deals with the design of provision schemes when the government wishes to supply certain private goods to the most needy groups in society. Ideally, the government would want to condition access to the public provision on an individual's underlying ability but only income is observable.

Thus, provision schemes must be either means-tested, at the risk of distorting income decisions, or universal. However, many publicly provided goods such as health care or education are consumed in discrete quantities, i.e., from a single supplier. As a result, a policy that makes a certain good available in a fixed quality to all citizens potentially distorts individual consumption choices.

In this chapter, I show that optimal provision schemes are de facto means-tested in the sense that the public supply is targeted at the poor. The optimum can be reached either through a formal means-test or under universal provision by subsidizing purchases of market alternatives to the publicly provided good. Intuitively, the government can avoid any income distortions from a means-test by adjusting the income tax to compensate those not included in the public program.

The result also has implications for the optimal tax treatment of private purchases of publicly provided goods in countries that have implemented universal provision of, e.g., health care. The chapter demonstrates that the optimal policy involves a subsidy that effectively decreases with income. The maximal subsidy is always less than the cost of servicing a single individual in the public sector.

Chapter 3, *Optimal Provision of Public Goods: A Synthesis*, is joint work with Claus Thustrup Kreiner. This chapter also deals with the government's role in goods provision but focuses on public goods, which by nature are enjoyed by everyone. Our analysis is motivated by the continuing contention that surrounds the academic literature on the optimal supply of public goods.

The existing literature is dominated by two different strands. One has no systematic link between expenditures and the financing scheme, besides budget balance, and emphasizes the distortionary costs of taxation and the marginal benefit of public goods as the main determinants of the optimal supply. The other applies the so-called benefit principle, which builds on the flexibility of the non-linear income tax and constructs the financing scheme to eliminate any distributional effects. This latter approach has focused on conditions when the original Samuelson rule applies.

In this chapter, we present a fully general analysis of the optimal supply of public goods based on the benefit principle. We derive an intuitive criterion for assessing the desirability of a marginal expansion of a public good that identifies the correlation between ability and the marginal willingness to pay for the public good as the determinant of any deviations from the Samuelson rule.

The chapter also offers a synthesis of the two different approaches in the literature by showing that both derive from the same basic formula. The only difference lies in the restrictions imposed on the financing scheme. In addition, we provide a generalization of the conditions for the applicability of the Samuelson rule. We show that correlations between income and the marginal willingness to pay that are not driven by ability do not lead to a departure from the first best. Finally, we demonstrate that the traditional focus on the MCF in analyses of public goods provision is only justified in special cases.

In Chapter 4, *Taxation, Regulation, and Cue-Triggered Consumption*, I study optimal

policy when demand is affected by external cues. Exposure to a visceral cue, as when a smoker sees a cigarette, increases consumption of the cue-related good. The major challenge for welfare economics is whether cue-induced demand effects reflect a rational response or misguidings of the mind.

The chapter presents a simple, unified framework that embodies rational choice as well as boundedly rational decisionmaking. The model permits an analysis not only of optimal tax policy but also of the effects of cue regulation. I show that it is optimal to subsidize cue-triggered consumption when there are no errors in individual decisionmaking. The optimal subsidy decreases with the degree of cue regulation. In contrast, it is optimal to tax sin good purchases when individuals are boundedly rational. The optimal sin tax decreases as cue regulation is tightened.

The fact that increased cue regulation leads to less reliance on the tax instrument is a novel result that holds regardless of the degree of individual rationality. I finally show that hassle policies such as smoking bans in public places are generally inefficient.

Chapter 5, *An Evaluation of the Tax-Transfer Treatment of Married Couples in European Countries*, is written jointly with Herwig Immervoll, Henrik Jacobsen Kleven, and Claus Thustrup Kreiner. We present a comprehensive analysis of various aspects of the combined tax and transfer treatment of married couples. The chapter uses the EU-ROMOD microsimulation model to carefully estimate individual tax burdens in the 15 pre-enlargement EU member states.

We begin our analysis by carefully mapping the precise nature of jointness in the tax-

transfer schemes in each of our sample countries. We find that many tax-transfer schemes in Europe display negative jointness where the tax rate on one person depends negatively on spousal earnings. The presence of negative jointness is driven by the interaction of family-based transfers and individual or almost-individual taxes. This finding also demonstrates that the usual sharp distinction between individual and joint tax systems is too simplified.

We next consider the incentives for secondary earners to supply labor and study the welfare effects of reforms that change the relative taxation of spouses. We set up a simple model of family labor supply and focus on the participation margin because the empirical findings indicate strong responsiveness of female labor supply along this margin along this margin. Our findings suggest that, in most European countries, lowering the tax burden on secondary earners is associated with large welfare gains.

Finally, the chapter simulates marriage penalties for hypothetical households. We find large marriage penalties at the bottom of the income distribution (but not at the top) in most countries. These are due primarily to the family-based transfers and therefore marriage penalties at the bottom in practice go hand in hand with negative jointness.

Summary in Danish

Denne afhandling består af fem kapitler, der kan læses uafhængigt af hinanden. Alle kapitler beskæftiger sig med statens rolle i samfundsøkonomien. Kapitel 1 omhandler efficiensomkostningerne forbundet med en progressiv skatteskala, når intertemporale responser tages med i betragtning. Kapitel 2 spørger om statslige ydelser bør være universelle eller målrettede. Kapitel 3 studerer det optimale udbud af offentlige goder, når staten er henvist til at benytte potentielt forvridende beskatning. Kapitel 4 anvender indsigter fra adfærdsøkonomi til at studere sammenhængen mellem beskatning og regulering af såkaldte *cue*-udløste forbrugsgoder. Endelig foretager Kapitel 5 en evaluering af skatte- og overførselssystemet rettet mod gifte par i Europa baseret på en stor mikrosimulation-model. Alle kapitler, bortset fra Kapitel 5, som udelukkende er anvendt, indeholder teoretiske bidrag. Kapitel 1 benytter endvidere amerikanske data til at drage en række politikkonklusioner.

Kapitel 1, *Intertemporal Tax Wedges and Marginal Deadweight Loss*, er skrevet sammen med Jes Winther Hansen. Dette kapitel kaster et nyt blik på effiçienoeffekterne af indkomstskattereformer i en eksplicit dynamisk model. Når indkomst beskattes årligt under en progressiv skatteskala vil den enkeltes marginalskat variere fra år til år. Dette resulterer i intertemporale skatteklir, der giver skatteyderen et incitament til at flytte indkomst til de år, hvor hans marginalskat er lavest.

Traditionelt har studier af dødvægtstabet forbundet med indkomstbeskatning benyttet en statisk ramme og dermed ignoreret intertemporal substitution, hvilket kan have

ført til fejlbehæftede estimater. I en fuldt dynamisk model formår vi at udtrykke det marginale dødvægtstab ved indkomstbeskatning alene som en funktion af empirisk observerbare elasticiteter og mængder. Udtrykket kan opdeles i en forvridning, der skyldes marginalskatte, og en forvridning, der kan tilskrives intertemporale skattekliver. Vi viser, at man ved at ignorere intertemporal substitution undervurderer det marginale dødvægtstab ved en skattereform, der øger de intertemporale skattekliver, og modsat overvurderer forvridningstab ved en reform, der mindsker de intertemporale skattekliver. I visse tilfælde kan en højere marginalskat endda føre til efficiensforbedringer, hvis den fører til en reduktion af de intertemporale skattekliver af tilpas størrelse.

For at få et indtryk af den kvantitative betydning af at ignorere intertemporal substitution i skattepligtig indkomst analyserer vi tre stiliserede skattereformer ved hjælp af et amerikansk paneldatasæt med omkring 34.000 repræsentative skatteenheder. Vi gennemfører simulationer for en bred vifte af elasticiteter for at afspejle den usikkerhed, der hersker i litteraturen omkring størrelsen af de intertemporale responser. Vi finder generelt, at det marginale dødvægtstab er særdeles følsomt overfor selv beskeden intertemporal substitution. Vores resultater antyder, at skatteprogressivitet er forbundet med større omkostninger end hidtil antaget.

Kapitel 2, *Public Provision of Private Goods: Optimal Provision Schemes*, beskæftiger sig med indretningen af den moderne velfærdsstat. Mere præcist studeres den mest hensigtsmæssige tilvejebringelse af private goder, som staten ønsker at stille til rådighed for de svageste grupper i samfundet. I en ideel verden ville staten betinge adgangen til

de offentlige ydelser på individual produktivitet, men kun inkomst er observerbar. Som et resultat må udbudet af offentlige goder enten målrettes ved hjælp af en indkomstgrænse (hvilket risikerer at forvråde indkomstbeslutningen) eller være universel. Mange offentligt tilvejebragte goder såsom sundhedsbehandling eller uddannelse forbruges dog kun i diskrete mængder fra en enkelt udbyder. Derfor risikerer en politik, der stiller et privat gode af en fastsat kvalitet til rådighed for alle borgere, at forvråde individuelle forbrugsbeslutninger.

Dette kapitel demonstrerer, at den optimale udformning af offentlige ydelser de facto anvender en indkomstgrænse til at målrette det offentlige udbud til de svageste grupper. Optimum kan nås enten ved brug af en formel indkomstgrænse eller med universelle ydelser ved at subsidiere private køb af alternativer til det offentligt tilvejebragte gode. Staten kan undgå, at indkomstgrænsen forvråder indkomstbeslutningen ved at tilpasse indkomstskatten, så de udeladte grupper bliver kompenseret.

Resultatet har også betydning for den optimale skattebehandling af private køb af ydelser, der i øvrigt tilvejebringes af det offentlige i et universelt system. Det kunne fx. være køb af private sundhedsforsikringer. Kapitlet viser, at den optimale politik inkluderer et subsidium, der aftager med indkomst. Det maksimale subsidium er altid mindre end omkostningen ved at servicere et enkelt individ i det offentlige system.

Kapitel 3, *Optimal Provision of Public Goods: A Synthesis*, er resultatet af et samarbejde med Claus Thustrup Kreiner. Dette kapitel beskæftiger sig ligeledes med statens rolle i tilvejebringelsen af visse ydelser, men fokuserer i stedet på offentlige goder, der kommer

alle til gode. Vores analyse er inspireret af den fortsatte uenighed, der hersker i den akademiske litteratur om det optimale udbud af offentlige goder.

Den eksisterende litteratur er domineret af to forskellige tilgange. Den første forudsætter ingen systematisk sammenhæng mellem udgifter til det offentlige gode og finansieringen, udover kravet om budgetbalance. Denne tilgang understreger forvriddingsomkostningerne ved beskatning og fordelingshensyn som de vigtigste faktorer, der bestemmer det optimal udbud. Den anden tilgang anvender det såkaldte *benefit principle*, der bygger på fleksibiliteten af den ikke-lineære indkomstskat og eliminerer enhver fordelingsmæssig effekt ved udformningen af finansieringsplanen. Denne anden tilgang har fokuseret på at identificere nødvendige betingelser, for at den oprindelige Samuelson regel gælder.

Kapitlet præsenterer en helt generel analyse af det optimale udbud af offentlige goder baseret på *benefit principle*. Vi udleder et intuitivt kriterie, der kan afgøre, hvornår en marginal udvidelse af udbudet af et offentligt gode er ønskværdig. Dette kriterie identificerer korrelationen mellem individuel produktivitet og betalingsvilligheden for det offentlige gode som afgørende for enhver afvigelse fra Samuelson reglen.

Kapitlet præsenterer samtidig en syntese af de to tilgange i den eksisterende litteratur ved at demonstrere, at de begge kan udledes fra den samme grundlæggende formel. Den eneste forskel består i de restriktioner, der pålægges finansieringsplanen. Vi præsenterer endvidere en generalisering af de nødvendige betingelser, for at Samuelson reglen holder. Det demonstreres, at korrelationer mellem indkomst og betalingsvillighed, der ikke skyldes den underliggende produktivitet, ikke fører til afvigelser fra Samuelson re-

glen. Endelig viser vi, at den traditionelle fokus på MCF i analyser af det optimale udbud af offentlige goder kun er berettiget i særlige tilfælde.

I Kapitel 4, *Taxation, Regulation, and Cue-Triggered Consumption*, studerer jeg den optimale politik overfor forbrugsgoder, hvor efterspørgslen er påvirket af såkaldte *cues*. Når en forbruger udsættes for et *cue*, som når en ryger ser en cigaret, øges forbruget af det relaterede gode. Den store udfordring for velfærdsøkonomi består i at afgøre, hvorvidt *cue*-påvirket efterspørgsel er udtryk for en rationel respons eller ej.

Kapitlet præsenterer en simpel ramme, der tillader både rationelle beslutninger og begrænset rationalitet. Modellen tillader en analyse af såvel den optimal skattepolitik som af effekterne af *cue* regulering. Jeg viser, at det er optimalt at subsidiere *cue*-udløst forbrug, når forbrugerne er fuldt rationelle. Det optimale subsidium er aftagende i graden af *cue*-regulering. Derimod er det optimalt at beskatte forbruget af syndgoder, når forbrugerne lider af begrænset rationalitet. Den optimale skat aftager med graden af *cue*-regulering.

Resultatet, at forøget *cue*-regulering fører til mindre anvendelse af skatteinstrumentet, er et nyt og gælder uanset graden af individuel rationalitet. Endelig demonstreres det, at politikker såsom forbud mod rygning på offentlige steder generelt er inefficente.

Kapitel 5, *An Evaluation of the Tax-Transfer Treatment of Married Couples in European Countries*, er skrevet i samarbejde med Herwig Immervoll, Henrik Jacobsen Kleven, og Claus Thustrup Kreiner. Vi præsenterer en omfattende analyse af forskellige aspekter af

skatte- og overførselssystemets behandling af gifte par. Kapitlet benytter EUROMOD mikrosimulationsmodellen til omhyggeligt at estimere individuelle skattebyrder i 15 EU-medlemslande (samtlige medlemslande fra før udvidelsen i 2004).

Vi indleder vores analyse med nøje at kortlægge den præcise karakter af afhængighed mellem ægtefæller i skatte- og overførselssystemerne i hvert af de 15 lande. Det viser sig, at mange skattesystemer i Europa udviser negativ afhængighed, hvor skattesatsen for den ene ægtefælle afhænger negativt af den anden ægtefælles indkomst. Tilstedeværelsen af negativ afhængighed skyldes samspillet mellem familie-baserede overførsler og individuel eller næsten-individuel beskatning. Dette mønster demonstrerer endvidere, at den traditionelle skarpe skelnen mellem individuelle og fælles skattesystemer er forsimplet.

Næste punkt i vores analyse er arbejdsudbudsincitamenterne for *second earners* og implikationerne heraf for velfærdseffekterne af forskellige reformer, der ændrer den relative beskatning af ægtefæller. Vi opstiller en simpel model for en families arbejdsudbud og fokuserer på deltagelsesbeslutningen, fordi empirien tyder på, at kvinders arbejdsmarkedsdeltagelse er særdeles følsom overfor skatteincitament. Vores resultater indikerer, at der i de fleste europæiske lande er store velfærdsgevinster forbundet med at reducere skattebyrden for *second earners*.

Endelig simulerer kapitlet *marriage penalties* for hypotetiske husholdninger. Vi finder, at der er store skattemæssige omkostninger ved ægteskab i bunden af indkomstfordelingen (men ikke i toppen) i de fleste lande. Disse kan primært tilskrives familie-baserede overførsler, og derfor går *marriage penalties* typisk hånd i hånd med negativ afhængighed.

Chapter 1

Intertemporal Tax Wedges and Marginal Deadweight Loss*

Jes Winther Hansen Nicolaj Verdelin
University of Copenhagen and EPRU

September 2007

Abstract

This paper demonstrates the importance of accounting for intertemporal substitution in taxable income when evaluating the efficiency loss of income taxation. Ignoring intertemporal substitution, which recent empirical evidence suggests is substantial, leads to biased estimates of the deadweight loss of taxation. This points to the problem of using a static approach to estimate the deadweight loss and holds significant implications for income tax progression. Using an empirically applicable expression for the marginal deadweight loss, we conduct simulations of tax reforms on U.S. panel data and find sizable biases.

Keywords: Intertemporal substitution; Intertemporal tax wedges; Marginal deadweight loss; Income taxation

JEL classification: H21; H24; D91

*We are grateful to Henrik Jacobsen Kleven, Claus Thustrup Kreiner, and seminar participants at the University of Copenhagen for helpful comments. All remaining errors are ours.

1 Introduction

Studies of income taxation have traditionally restricted attention to static models. In practice, however, income is taxed annually, which makes the study of income taxation an inherently dynamic question. The combination of annual income taxation and a progressive tax schedule causes individual marginal tax rates to change when the taxpayer's income varies over time. This recurrent variation in marginal tax rates gives rise to *intertemporal tax wedges* between different years. These wedges distort the intertemporal allocation of income by creating an incentive for intertemporal substitution because taxpayers can reduce their overall tax liability by shifting income to years when they face low marginal tax rates. This paper demonstrates that intertemporal substitution in taxable income is important for assessing the efficiency of tax reforms — even for modest elasticities of intertemporal substitution. The previous studies by Harberger (1964), Browning (1987), and Feldstein (1999) that have developed and used the static framework to evaluate the deadweight loss of income taxation have ignored intertemporal substitution and may have produced biased estimates.

We simulate marginal tax reforms on U.S. data for a range of values for the intertemporal elasticity of substitution in order to quantify the biases from intertemporal responses. To lay the foundation for the empirical simulations, we first derive a simple, empirically applicable expression for the marginal deadweight loss of income taxation in a dynamic setting. We show that the distortions from a marginal tax reform arises both from the level of the marginal tax rate and from the intertemporal tax wedges. Account-

	<i>Percent of group</i>	<i>Average MTR</i>	<i>Average absolute MTR change (pct. points)</i>
<i>Full sample</i>		37.7%	7.6
<i>At least one tax change</i>	76.1%	38.6%	10.0
<i>One change</i>	25.1%	37.9%	10.6
<i>Two changes</i>	26.6%	38.4%	10.5
<i>Three changes</i>	24.4%	39.5%	8.7
<i>High-income taxpayers</i>			
<i>At least one tax change</i>	93.2%	45.4%	5.5

Table 1: Intertemporal tax wedges, 1996–1999.

The first column shows the distribution of taxpayers according to the number of times their marginal tax rate changed between 1996 and 1999. High-income taxpayers have taxable income above \$100,000. The second column lists the average marginal tax rates over the period for these groups. The third column displays the average absolute difference in marginal tax rates between contiguous years. The table is based on a trimmed sample that, e.g., excludes observations with a change in marital status (see Section 4). Sources: Survey of Income and Program Participation, 1996 panel, and NBER’s TAXSIM model.

ing for intertemporal substitution leads to higher estimates of the marginal deadweight loss if the reform increases intertemporal tax wedges, e.g., by raising the tax rate in one of the top brackets. Analogously, it leads to lower estimates if intertemporal tax wedges are reduced. In some cases, a marginal tax increase can even improve overall efficiency if the reduction in intertemporal tax wedges is sufficiently strong. The underlying intuition resembles the Ramsey tax problem where the efficiency cost of a tax on one good (income in one year) depends on pre-existing taxes on other goods (income in other years). These insights are important for evaluating the costs of, e.g., the Earned Income Tax Credit, time-limited transfers such as the U.S. dependent exemption, and the progression of income tax rates.

The practical importance of intertemporal tax wedges can be illustrated by considering

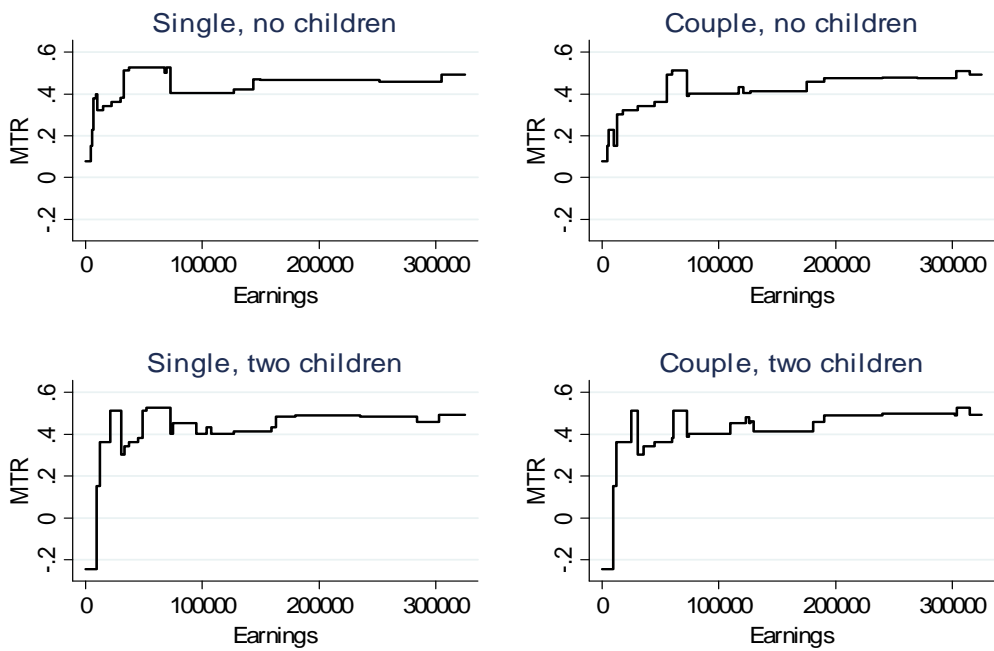


Figure 1: Effective marginal tax rates in California for different tax units, 1999. Combined state and federal effective marginal income tax rates. Source: NBER's TAXSIM model.

the nationally representative Survey of Income and Program Participation (SIPP) sample for the U.S. Table 1 illustrates the year-to-year variation in marginal tax rates for this sample. The first column shows the distribution of taxpayers according to the number of times their marginal tax rate changed between 1996 and 1999. An estimated 76% percent of survey respondents experienced at least one change and as many as 24% saw changes every year. The second column lists the average marginal tax rates over the period for each of these groups. The final column illustrates the size of the variation in marginal rates by calculating the average absolute difference in marginal tax rates between contiguous years. For instance, the marginal tax rate changed on average 8.7 percentage points

from year to year for the group of taxpayers experiencing three changes. Changes in the marginal tax rate are even more common for high-income taxpayers. Among those taxpayers with taxable income in excess of \$100,000 in at least one year, a full 93% experienced a marginal tax change. Overall, the table shows that intertemporal tax wedges are common and sizable. The presence of intertemporal tax wedges reflects the complex shape of the U.S. income tax schedule as illustrated in Figure 1, which displays marginal tax rates for different families with taxable income below \$325,000. The kinks in the tax schedule imply that marginal tax rates may change considerably from year to year, even for small changes in income.

Taxpayers can substitute taxable income intertemporally through an intertemporal reallocation of labor supply or by merely changing the timing of income. Labor supply can be adjusted by working overtime or taking on additional jobs, the timing of salaries and bonuses is flexible to some degree, and there is almost full discretion over the exercise of stock options and the realization of asset gains. Generally, high-income earners are expected to have more flexibility in adjusting their incomes but many of these options are also available to lower income groups. The opportunity to take advantage of intertemporal tax incentives depends crucially on taxpayers' ability to anticipate future marginal tax rates. Some events are difficult to predict, e.g., divorces, layoffs, and loss of earnings ability due to illness. Others are more easily anticipated such as promotions, tax bracket creep, and expiration of time-limited tax credits. All these factors contribute to the individual propensity to substitute income over time.

While intertemporal substitution has long been recognized in the empirical literature on labor supply, only a few studies have produced direct estimates of elasticities of intertemporal substitution in taxable income. This response has commonly been regarded as a short-term, temporary adjustment around the time of a tax reform. It is important to recognize, however, that an annual progressive income tax gives a recurrent incentive to shift income intertemporally. Goolsbee (2000) studies the OBRA93 tax reform, which raised marginal tax rates at the top. He finds evidence of substantial intertemporal substitution for high-income earners around the time when the reform was implemented. While Goolsbee focused on a tax reform, Looney and Singhal (2006) exploit the expiration of the U.S. tax exemption for dependents to estimate the response to year-to-year variation in marginal tax rates. Their findings suggest considerable intertemporal shifting of earned income for families with middle incomes around \$60,000. Overall, the limited empirical evidence indicates that intertemporal substitution in taxable income is non-negligible.

In our empirical simulations, we consider three tax reforms for the U.S. First, a marginal increase in the 36% federal rate, which was the second-highest federal rate applicable in 1999. We find that accounting for intertemporal substitution substantially increases the estimate of the marginal deadweight loss of the 36% rate in accordance with the prediction that the reform increases intertemporal tax wedges. Second, we consider an increase in the lowest federal rate, the 15% rate. As expected, accounting for intertemporal responses decreases the estimated marginal deadweight loss of the 15% rate. Finally, we analyze a more complicated reform that introduces phase-out of the

personal tax exemption. Because it lowers a severe intertemporal tax wedge, this reform improves efficiency in many scenarios. As a preview of our numerical results, we estimate the marginal deadweight loss of the 15% federal rate in a static framework to be 7% of tax revenue for an elasticity of taxable income of 0.2. Accounting for tax bracket mobility and intertemporal substitution, with an intertemporal cross-price elasticity of -0.05 , the estimate drops to 5.4%, which is a decrease of 23%. A cross-price elasticity of -0.05 corresponds to, e.g., a mere 0.5% decrease in taxable income this year in response to a 10% increase in the net-of-tax marginal rate next year. The sensitivity of the estimated marginal deadweight loss to intertemporal responses, even when these are very small, is a general finding of our empirical simulations.

The study of income taxation in a dynamic setting has recently attracted considerable interest. Werning (2007) demonstrates the optimality of marginal tax smoothing with a non-linear tax in a model with aggregate uncertainty when there is no idiosyncratic skill mobility. In another recent paper, Gaube (2007) solves for the optimal income tax in a two-period, two-type model under the assumption that income is taxed annually. While we do not attempt to solve the optimal dynamic tax problem, we provide an empirical account of the intertemporal efficiency costs that are a key determinant of optimal taxes.

The paper proceeds as follows. In Section 2, we set up a dynamic model of income taxation and present our theoretical analysis. We begin the quantitative analysis in Section 3 by considering the problem of calibrating behavioral elasticities and proceed to the empirical simulations in Section 4. Section 5 concludes.

2 Marginal Deadweight Loss in a Dynamic Setting

By far the most prevalent way of taxing income is by means of an annual income tax. Although this can be supplemented by some redistribution based on lifetime income (e.g., asset-tested benefits for retirees), such policies constitute a minor part of most tax systems. Annual income taxation gives taxpayers an incentive to shift income over time in order to reduce their tax liabilities when facing a progressive tax schedule.

In this section, we develop an expression for the marginal deadweight loss of income taxation in a dynamic setting. This expression allows us identify the bias from ignoring intertemporal substitution in taxable income and forms the basis for the empirical simulations in Section 4. The novel aspect of our approach is to isolate the distortions from the intertemporal tax wedges. These dynamic distortions are precluded in previous studies using a static approach.

The analysis builds on a simple extension of a static model with two goods. Each taxpayer has a finite planning horizon of N periods and a well-behaved utility function $U(c_1, z_1, \dots, c_n, z_n, \dots, c_N, z_N)$, where the subscript n refers to the time index. The taxpayer gets utility from a consumption good, c , and disutility from income, z . To simplify the exposition the interest rate is assumed to be zero, but the model can readily be generalized to incorporate discounting. Also, there is no uncertainty and no restrictions on savings between periods. Income is taxed in each period using a piece-wise linear tax schedule. The tax function $T(z)$ constitutes a net payment to the public sector, embodying both taxes and transfers, and is constant over time. It is convenient to express the taxpayer's

problem as an expenditure minimization problem

$$\min_{\{c_n, z_n\}} \sum_{n=1}^N [c_n - z_n + T(z_n)] - \lambda [U(c_1, z_1, \dots, c_N, z_N) - \bar{u}], \quad (1)$$

where λ is a Lagrange multiplier and \bar{u} is the utility level. From (1) we obtain a sequence of compensated demand functions, $c_n = c_n(T, \bar{u})$, and compensated income supply functions, $z_n = z_n(T, \bar{u})$.¹ The value function for the minimization problem is the expenditure function, $e(T, \bar{u}) = \sum [c_n - z_n + T(z_n)]$.

The deadweight loss of taxation (also called the “excess burden”) is based on a conceptual experiment where the government imposes taxes, thereby distorting prices, and returns the revenue to the taxpayer lump sum. The deadweight loss is the amount of income in excess of tax revenue that the taxpayer is willing to give up in return for a removal of all taxes.² The efficiency impact of a tax reform is expressed as the change in the deadweight loss, which is the change in taxpayer expenditure less the change in tax revenue

$$\Delta DWL = e(\hat{T}, \bar{u}) - e(T, \bar{u}) - \sum_{n=1}^N [\hat{T}(\hat{z}_n) - T(z_n)],$$

where $T(z)$ is the pre-reform tax function and $\hat{T}(\hat{z})$ is the post-reform tax function. The associated sequence of compensated income supply functions are z_n and \hat{z}_n , respectively.

In this section, we consider a tax reform consisting of a change in the tax rate t^j in a single tax bracket j .³ The change in t^j changes the tax function $T(z)$ in all periods.

¹We do not consider taxpayers who are located at kink points in the tax schedule. Theoretically, there may be bunching at kink points, but the literature has not found evidence that such behavior is widespread, e.g., Saez (2002). Further, if the income distribution is continuous and there are a finite number of kink points, individuals located at these points have measure zero.

²See Auerbach (1985) for a thorough theoretical exposition of the deadweight loss of taxation.

³In Appendix B, we analyze a general tax reform where marginal tax rates in several tax brackets

However, the taxpayer only experiences a marginal tax change in periods when she is taxed in bracket j at the margin. Since the reform changes $T(z)$ in all periods, it can be interpreted as a permanent tax reform before the first period. This is similar to the reform underlying the deadweight loss in the static model, except that in our case the single static period is partitioned into N sub-periods.⁴

The marginal deadweight loss of t^j for each taxpayer is the efficiency effect of a marginal change in t^j . We use the following notation: τ_n is the taxpayer's marginal tax rate in period n and Ω is the set of periods where $\tau_n = t^j$, i.e., the periods where the taxpayer's marginal tax rate changes (because the taxpayer is taxed in bracket j at the margin). The intertemporal tax wedge between periods n and m is the absolute difference in marginal tax rates, $|\tau_n - \tau_m|$. We have

$$\begin{aligned} \frac{dDWL}{dt^j} &= \frac{de(T, \bar{u})}{dt^j} - \sum_{n=1}^N \frac{dT(z_n)}{dt^j} \\ &= \sum_{n=1}^N \left\{ \frac{\partial T(z_n)}{\partial t^j} - \sum_{m \in \Omega} \left[\frac{\partial c_n}{\partial (1 - \tau_m)} - (1 - \tau_n) \frac{\partial z_n}{\partial (1 - \tau_m)} \right] \right\} \\ &\quad - \sum_{n=1}^N \left[\frac{\partial T(z_n)}{\partial t^j} - \tau_n \sum_{m \in \Omega} \frac{\partial z_n}{\partial (1 - \tau_m)} \right]. \end{aligned}$$

In each period there are effects on expenditure from the mechanical change in tax liability (absent behavioral responses) and from responses in c and z to the marginal tax rate changes in all periods $m \in \Omega$. Analogously, tax revenue in each period changes both

change simultaneously.

⁴Alternatively, the deadweight loss can be interpreted as the efficiency cost of a tax reform from the moment the reform is announced. This requires that we regard bracket j before and after the reform as two separate tax brackets. However, the model cannot capture the deadweight loss from unexpected changes in the tax schedule.

mechanically and as a result of responses in taxable income.

By the envelope theorem, behavioral responses have no first order effects on the taxpayer's minimized expenditure. As a result, the marginal deadweight loss can be expressed solely as the revenue implications of the responses in taxable income

$$\frac{dDWL}{dt^j} = \sum_{n=1}^N \frac{\tau_n}{1-t^j} z_n \sum_{m \in \Omega} \varepsilon_{nm}, \quad (2)$$

where $\varepsilon_{nm} \equiv [(1 - \tau_m) / z_n] \cdot \partial z_n / \partial (1 - \tau_m)$ is the compensated elasticity of taxable income in period n with respect to the marginal net-of-tax rate in period m .⁵ We have also used that $\tau_m = t^j$ for all $m \in \Omega$ by definition. The compensated elasticity of taxable income neatly sums up all relevant behavioral responses, including labor supply, tax avoidance, and the form of compensation, as emphasized by Feldstein (1999), as well as changes to the timing of income, whether due to intertemporal substitution in labor supply or the timing of reporting.

For each period where t^j is the marginal tax rate, there (a) is an *intratemporal* distortion of income within the period, captured by the own-price elasticity, ε_{mm} , and (b) are changes in income in all other periods because of *intertemporal substitution*, which is captured by the cross-price elasticities, ε_{nm} for $n \neq m$. Compensated own-price elasticities are always non-negative. The intertemporal cross-price elasticities are intuitively expected to be non-positive such that an increase in τ_m implies a (partial) increase in period n income, which is now relatively cheaper. Throughout this paper we will assume

⁵Only compensated responses count towards the marginal deadweight loss. Income effects do not matter since the definition of the deadweight loss assumes that tax revenue can be returned to the taxpayer lump-sum.

that the cross-price elasticities are negative, but at this point we do not impose further restrictions. Expression (2) is easily empirically applicable since it is based on elasticities that are, in principle, directly observable. Crucially, these elasticities are not structural parameters but depend on the economic environment, the tax schedule in place, and the tax reform in question.

The ultimate aim of this section is to provide a framework allowing us to assess the bias from ignoring intertemporal substitution in taxable income when computing the marginal deadweight loss. In order to achieve this we introduce the *static elasticity*, η . The static elasticity is defined as the compensated response in income to a permanent change in the marginal net-of-tax rate for a taxpayer who faces the same marginal tax rate, τ , in all periods. Since η is derived for a constant marginal tax rate, the static response involves no intertemporal substitution in taxable income due to intertemporal tax wedges

$$\eta_n \equiv \frac{1 - \tau}{z_n} \frac{dz_n}{d(1 - \tau)} = \sum_{m=1}^N \varepsilon_{nm}. \quad (3)$$

This response is equivalent to the response in a static model, which, by definition, has no tax bracket mobility. Using (3), we can rewrite (2) as (see Appendix A for details)

$$\begin{aligned} \frac{dDWL}{dt^j} &= \frac{t^j}{1 - t^j} \sum_{m \in \Omega} z_m \eta_m \\ &+ \sum_{m \in \Omega} \sum_{n \neq m} \left(\frac{\tau_n}{1 - t^j} z_n \varepsilon_{nm} - \frac{t^j}{1 - t^j} z_m \varepsilon_{mn} \right). \end{aligned} \quad (4)$$

Equation (4) identifies the bias from ignoring intertemporal tax wedges and substitution in taxable income. The first term is the *static effect*. This reflects that an increase in

the marginal tax rate, t^j , creates an efficiency loss because the taxpayer substitutes away from income toward leisure. This distortion is present in all periods where the taxpayer faces the marginal tax increase and arises from the level of t^j . The static effect disregards tax bracket mobility and hence corresponds to the traditional static Harberger-Browning formula for the marginal deadweight loss. In addition, there is a *dynamic effect* captured by the second term. Individual marginal tax rates vary over time due to the shape of the tax schedule. When t^j increases, income is substituted over time toward other tax brackets, which are now relatively cheaper. As a consequence, there is a revenue loss from substitution toward periods with lower marginal tax rates and a revenue gain from substitution toward periods with higher marginal tax rates. This is a distortion from intertemporal substitution in taxable income.

In Section 4 we evaluate (4) on a panel data set in order to quantify the bias. Failing to take account of intertemporal substitution, as the static approach does, corresponds to assuming that $\varepsilon_{nm} = 0$ for $n \neq m$. Although ignoring a margin of behavioral response always leads to an underestimate of the total deadweight loss, the marginal deadweight loss can be either positively or negatively biased depending on how the tax reform affects intertemporal tax wedges. From the dynamic effect in (4), the bias tends to be negative if intertemporal tax wedges are increased. In this case, the taxpayer generally faces marginal tax rates that are lower than t^j and hence raising t^j increases the distortion from intertemporal substitution. Likewise, the bias tends to be positive if intertemporal

tax wedges are reduced.⁶ In the latter case, there is the possibility that a marginal tax increase can improve overall efficiency. The simple second-best explanation is that the intratemporal distortion associated with a tax increase can be dominated by efficiency gains from reductions of intertemporal distortions. The reverse can also be true: lowering a marginal tax rate can reduce efficiency if intertemporal tax distortions are exacerbated.

Results:

1. *Ignoring intertemporal substitution in income underestimates the marginal deadweight loss of income taxation when intertemporal tax wedges increase.*
2. *Ignoring intertemporal substitution in income overestimates the marginal deadweight loss of income taxation when intertemporal tax wedges decrease.*
3. *In some cases, an increase in a positive marginal tax rate can improve efficiency if intertemporal tax wedges decrease.*

On a general note, (4) reveals that intertemporal substitution in taxable income adds to the efficiency costs of a progressive income tax. When there is income mobility, tax progression creates intertemporal tax wedges that distort the timing of income. Importantly, income mobility in general also weakens the redistributive motive for tax progression by making the distribution of lifetime income less unequal than the distribution in a

⁶If marginal tax rates are progressive, such that $\tau_m \geq \tau_n \iff z_m \geq z_n$, there is a clear case for the biases described in the text. The biases can then only be of the opposite sign if there are strong asymmetries in the cross-price elasticities.

single year. Although our analysis does not pretend to speak of optimality, these insights suggest that progressive income taxation may be less desirable than previously thought.

3 Empirical Evidence on Behavioral Elasticities

The analysis in Section 2 demonstrated that there can be biases from ignoring intertemporal substitution in taxable income when calculating the marginal deadweight loss of taxation. In the remainder of the paper we simulate tax reforms using U.S. data in order to stress the quantitative importance of these biases. The simulations require data on tax bracket mobility, i.e., information on income and marginal tax rates for a panel of taxpayers. We also need estimates of the behavioral elasticities in order to calibrate the model. In this section, we briefly review and discuss the relevant empirical evidence on the elasticity of taxable income.

A key question is the size of the intertemporal cross-price elasticities of taxable income but only a small number of empirical studies have produced direct estimates of these elasticities. One common approach to estimating the behavioral responses to taxation is to use a tax reform as a source of exogenous variation in marginal tax rates. Most studies using this methodology attempt to eliminate the influence of short-run intertemporal substitution on the estimation in order to obtain estimates of the permanent response. The reason is that the short-run substitution is considered a temporary response around the time of the reform. However, as stressed above, the incentive to shift income intertemporally is a recurrent phenomenon when there is tax bracket mobility and the short-run

responses may very well be important. Goolsbee (2000) finds evidence of substantial short-run intertemporal substitution in response to the Omnibus Budget Reconciliation Act in 1993 (OBRA93) among high income corporate executives. The main specification results in an elasticity of taxable income with respect to the current net-of-tax rate above one and an elasticity with respect to the net-of-tax rate the following year of about -0.8 . This suggests that most of the response to OBRA93 was transitory and that the permanent response has an elasticity of less than 0.4. Importantly, most of the response appears to be driven by the exercise of stock options. The corresponding elasticities for salaries and bonuses only are much smaller and the cross-price elasticity is statistically insignificant. It is difficult to generalize the results, both because of the special nature of executive compensation and because it is likely that OBRA93 was only anticipated in the last few months of 1992, leaving a short time for anticipatory responses.

Instead of using direct estimates, the compensated intertemporal cross-price elasticities can be approximated by exploiting the different responses to expected and unexpected marginal tax changes. The relationship between any two periods $m \neq n$ on the intertemporal income path, assuming that there are no *unexpected* changes in tax rates, can be approximated by

$$\ln z_n = \ln z_m + \gamma_{nm} [\ln(1 - \tau_n) - \ln(1 - \tau_m)] + \Delta_{nm}, \quad (5)$$

where γ_{nm} is the Frisch elasticity and Δ_{nm} is an error term assumed independent of the tax schedule.⁷ The Frisch elasticity measures the response in income to *expected* changes

⁷Empirically, factors that influence income can also affect marginal tax rates directly, implying that

in the net-of-tax marginal rate along the income path.⁸ A tax reform, in effect from the first period onwards, that changes the net-of-tax marginal rate in period m causes a reallocation of income such that the entire intertemporal path of income shifts. The compensated intertemporal cross-price elasticity of income in period n with respect to $1 - \tau_m$ measures how much z_n shifts as part of this response. Since (5) also holds on the new income path, we can find the compensated intertemporal elasticity by differentiating (5) with respect to $1 - \tau_m$

$$\varepsilon_{nm} = \varepsilon_{mm} - \gamma_{nm}, \quad (6)$$

where $\varepsilon_{nm} \equiv [(1 - \tau_m) / z_n] \partial z_n / \partial (1 - \tau_m)$ is the compensated elasticity as before. It can be shown theoretically that $\gamma_{nm} \geq \varepsilon_{mm}$ for all m (see, e.g., MaCurdy, 1981). Hence, the intertemporal cross-price elasticity is non-positive: an increase in the marginal tax rate in period m increases income in period n because income is substituted toward the relatively cheaper period n . We assume that both ε and γ are constant over time, such that we can use (6) to parameterize our simulations using only the static elasticity η and the Frisch elasticity γ . In this case, η is also constant and non-negative and from (3) the cross-price elasticities are $\varepsilon_{nm} = [\eta + (N - 1)\gamma] / N - \gamma$ for all m and n . The assumption of constant elasticities is not necessarily very realistic since it implies that ε_{nm} is independent of the time interval between period m and period n . Intertemporal

the error term could be correlated with the tax rates. An example is the birth of a child. Our approximation is valid as long as consistent estimates of the Frisch elasticity γ_{nm} are available, i.e., estimates that capture the correct, partial effect of marginal tax rates on income.

⁸Mathematically, the Frisch elasticity is the response in earned income to changes in the net-of-tax rate for a fixed marginal utility of income. The Frisch elasticity is often referred to as the intertemporal elasticity of substitution.

cross-price responses are likely lower for longer time horizons, e.g., because of uncertainty about future income, preferences, etc. Nevertheless, because there is not much empirical evidence on this relationship, our main simulations will be calibrated assuming constant cross-price elasticities. We conduct a robustness check of our results by relaxing this assumption at the end of Section 4.

There is still considerable uncertainty in the empirical literature about the absolute and relative magnitudes of these elasticities. Whether the Frisch elasticity or the compensated elasticity is estimated depends on whether the variation in marginal tax rates is expected or unexpected and most studies do not explicitly consider this aspect. In a seminal paper, MaCurdy (1981) puts forward a framework for estimating intertemporal labor supply and emphasizes the difference between expected and unexpected changes in the net wage. Estimating on a panel of white, married males, he finds Frisch elasticities for labor supply in the range 0.1–0.2 and very small cross-price elasticities. Feldstein (1995) exploits a panel of tax returns to study the elasticity of taxable income. He uses the Tax Reform Act of 1986 as a natural experiment and finds sizable elasticities well above one for high income taxpayers. Gruber and Saez (2002) employ a similar empirical strategy but use variation in marginal tax rates from various U.S. tax reforms in the 1980's and add elaborate controls for mean reversion and distributional changes.⁹

Their main specification yields an estimate of 0.4 for taxable income and a somewhat

⁹Interestingly, mean reversion in the U.S. income distribution is well-documented in the literature on the elasticity of taxable income. This supports the importance of a dynamic setting for the study of income taxation.

lower estimate for broad income. The results are mainly driven by responses among high income taxpayers. Kopczuk (2005) reviews the earlier literature and demonstrates that the results are sensitive to econometric specifications and sample selection effects. Kopczuk also stresses that the availability of deductions is a very important determinant of the elasticity of taxable income. In a recent paper, Ljung and Ragan (2006) estimate the response in earned income to a tax reform in Sweden. Their sample is a large and detailed panel of administrative data and they use both static and dynamic empirical specifications. Their preferred estimate is 0.37 for the compensated elasticity of earned income. In another interesting recent paper, Looney and Singhal (2006) estimate the responses in income to changes in marginal tax rates due to the expiration of the tax exemption for dependents, which occurs when the dependent turns 19 years old. Looney and Singhal argue that any resulting changes in marginal tax rates have been expected well in advance such that their estimate can be interpreted as a Frisch elasticity. They find an elasticity of earned income of 0.75 for middle incomes.

Based on the above evidence, we find that a reasonable range for η and γ is 0.2–0.8 for high incomes and somewhat smaller for low incomes. As a result, we consider cross-price elasticities in the intervals -0.05 to -0.4 for high incomes and -0.01 to -0.1 for low incomes reasonable. These ranges are, we believe, quite conservative. The reason is that we wish to demonstrate the importance of recognizing intertemporal income shifting, even if the responses are small.

4 Simulated Tax Reforms on U.S. Data

4.1 Data

Our data set is from the 1996 panel of the U.S. Survey of Income and Program Participation (SIPP), which covers the years 1996 through 1999. SIPP is a nationally representative sample and the 1996 panel consists of 116,000 individuals. The data is collected through interviews every four months with an appointed reference person in each household, who provides information for all its members. The survey contains information on earnings and other sources of income ranging from interest income to retirement pensions, demographics including age and gender, and family relationships. The latter allow us to match individuals in the sample to construct families and couples. In addition to information about different sources of income, our tax rate calculations require knowledge of filer status and the number of qualifying dependents, including those eligible for the Child Tax Credit (CTC). We identify dependents as those sample members who are younger than 19, or 24 if they are full time students, and count them as eligible for the CTC in a given year if they are under the age of 17 by the end of that year. Further, we maintain the assumption that married couples file jointly and that single individuals file as head of household whenever they have dependents. We exclude observations with a change in marital status because the resulting tax implications are large and not likely to be anticipated. We do not have information on itemized deductions and assume that everyone claims the standard deduction. After trimming the sample by excluding tax units with missing earnings data at some point during the panel, we end up with 33,826

tax units.

To eliminate the influence on marginal tax rates from unexpected changes in legislation, we convert all income figures into 1999-dollars and use the 1999 tax legislation in all years.¹⁰ We then compute effective marginal tax rates, not accounting for benefits, but including the Earned Income Tax Credit (EITC) and the Alternative Minimum Tax, using NBER's TAXSIM model.¹¹ Our calculations assume that the EITC take-up rate is 100% among eligible taxpayers. Although it is not part of our main simulations, we also use information in SIPP on receipt of TANF, Food Stamps, and SSI, combined with state and federal rules, to compute phase-out rates for these programs. This allows us to calculate effective marginal tax rates including benefit phase-out.

4.2 Simulations

In this section, we present the results of three simulated tax reforms. The purpose of our simulations is to shed light on the quantitative importance of the distortions caused by intertemporal tax wedges. As emphasized throughout the paper, these efficiency costs would not be captured in a static framework. We calculate the marginal deadweight loss using equation (B-2) from Appendix B. This formula is similar to (4) but allows for simultaneous changes in multiple tax brackets. The generalized formula is necessary because the marginal tax rates may change between years for taxpayers who are affected

¹⁰Since there were no federal tax reforms between 1996 and 1999, the unexpected variation could only come from changes to state tax legislation.

¹¹NBER's TAXSIM model is available on the Internet: www.nber.org/taxsim.

by a given reform in more than one year.¹² We use the panel weights included in SIPP and sum the individual marginal deadweight losses to obtain estimates that are representative of the U.S. population. Throughout, the marginal deadweight loss is expressed as a percentage of total tax revenue.

We present the results of our simulations for different values of the static elasticity η and the cross-price elasticity ε_{nm} . This sensitivity test reflects the uncertainty in the literature about the central elasticities. Further, we consider different ranges for the two elasticities across reforms to reflect the composition of the group of affected taxpayers. For each set of parameter values, we express the bias from ignoring intertemporal substitution by computing the percentage deviation of the marginal deadweight loss from the estimate obtained for the same value of η under the assumption that all cross-price elasticities are zero (the values in the first row of each table).

Reform 1: Marginal tax increase in the 36% federal tax bracket

The first reform raises the marginal tax rate for all taxpayers who are liable, at the margin, for the 36% federal rate, which was the second-highest federal marginal rate applicable in 1999. Specifically, we raise the tax rate for taxpayers reporting taxable income in the range \$130,250–\$283,150 (single filers), \$144,400–\$283,150 (head of household filers), and \$158,550–\$283,150 (married, filing jointly). On average, 313 tax units in our sample are affected by the reform in each of the four years, corresponding to 2,259,908 individuals

¹²We also correct for the difference in tax bases that apply to the federal tax schedule and earnings-based taxes and transfers. For instance, using only taxable income as the tax base would not capture the marginal deadweight loss of the EITC for taxpayers with zero taxable income but positive earnings.

Cross-price elasticity ε_{nm}		Static elasticity η				
		0.2	0.3	0.4	0.6	0.8
0	DWL:	3.57	5.36	7.14	10.70	14.30
	%-dev.:	0	0	0	0	0
-0.05		4.14	5.93	7.72	11.29	14.86
		16.04	10.69	8.02	5.35	4.01
-0.1		4.72	6.50	8.29	11.86	15.43
		32.08	21.39	16.04	10.69	8.02
-0.2		5.86	7.65	9.44	13.01	16.58
		64.17	42.78	32.08	21.39	16.04
-0.4		8.16	9.94	11.73	15.30	18.87
		128.34	85.56	64.17	42.78	32.08

Table 2: Change in the 36% federal bracket (Reform 1).

each year using panel weights. This reform affects only high income filers, approximately the top 2% of the taxable income distribution. On average, the reform is expected to increase the intertemporal tax wedges since many of the affected individuals are likely to face lower marginal tax rates in adjoining years if their marginal tax bracket changes.

We consider values for η of 0.2, 0.3, 0.4, 0.6, and 0.8, and for ε_{nm} of 0, -0.05 , -0.1 , -0.2 , and -0.4 . These values are motivated by the fact that the literature has demonstrated fairly high responsiveness for high income earners. Results of the simulation exercise for Reform 1 are in Table 2.

In general, the estimates for the marginal deadweight loss are sizable because of the high average income of the affected taxpayers. Not accounting for intertemporal

substitution leads to substantial underestimates of the marginal deadweight loss of the 36% federal tax rate. This confirms the intuition that a number of the affected individuals face a lower marginal tax rate at some point during the panel years. As an example, if $\eta = 0.4$ the baseline static marginal deadweight loss, shown in the first row of the table, is estimated to be 7.14% of tax revenue. This estimate ignores intertemporal substitution in taxable income and corresponds to the static effect. If instead the cross-price elasticity is assumed to be -0.1 , the estimate changes to 8.29% of tax revenue, which is an increase of 16%. The additional efficiency loss reflects that a number of individuals affected by the reform shift income to years in which they are not subjected to the tax increase.

Reform 2: Marginal tax increase in the 15% federal tax bracket

The second reform involves an increase in the marginal tax rate for all taxpayers in the 15% federal tax bracket. This corresponds to individuals reporting taxable income in the range \$0–\$25,750 (single filers), \$0–\$34,550 (head of household filers), and \$0–\$43,050 (married, filing jointly). There is a high density of taxpayers in these income ranges and on average the reform affects 24,507 tax units, or 72% of our sample, in each of the four years. This corresponds to 93,633,068 individuals using panel weights.

We choose lower values for the elasticities than those used for the high income earners in Reform 1 to reflect the findings in the literature that low income taxpayers are less responsive. Specifically, we consider values for η of 0.1, 0.15, 0.2, 0.3, and 0.4, and for ε_{nm} of 0, -0.01 , -0.02 , -0.05 , and -0.1 . The simulation results are in Table 3.

Ignoring intertemporal substitution implies large overestimates of the marginal dead-

Cross-price elasticity ε_{nm}		Static elasticity η				
		0.1	0.15	0.2	0.3	0.4
0	DWL:	3.52	5.27	7.03	10.55	14.06
	%-dev.:	0	0	0	0	0
-0.01		3.19	4.95	6.71	10.22	13.74
		-9.23	-6.15	-4.62	-3.08	-2.31
-0.02		2.87	4.62	6.38	9.90	13.41
		-18.46	-12.31	-9.23	-6.15	-4.62
-0.05		1.89	3.65	5.41	8.92	12.44
		-46.15	-30.77	-23.08	-15.38	-11.54
-0.1		0.27	2.03	3.79	7.30	10.82
		-92.31	-61.54	-46.15	-30.77	-23.08

Table 3: Change in the 15% federal bracket (Reform 2).

weight loss. This reflects that the 15% rate applies to the lowest federal tax bracket, which implies that taxpayers are likely to face higher marginal tax rates in surrounding years. Raising the 15% rate is thus likely to reduce intertemporal tax wedges. As an example, the static marginal deadweight loss is estimated to be 7.03% of tax revenue if $\eta = 0.2$. If we allow for intertemporal substitution corresponding to, e.g., $\varepsilon_{nm} = -0.05$, the estimate changes to 5.41% of tax revenue, which is a decrease of 23%. Because this reform affects such a large fraction of taxpayers, the estimates of the marginal deadweight losses are sizable such that even a small bias is quantitatively important.

Reform 3: Marginal phase-out of exemptions

This reform raises the marginal tax rate for taxpayers reporting zero taxable income but

who have positive AGI, and who did not receive benefits. We interpret this reform as a marginal phase-out of personal tax exemptions, i.e., a tax on earnings for individuals reporting income below the threshold for the 15% federal rate. In order to keep benefit recipients unaffected by the reform, benefit phase-out rates would have to be reduced in the same interval. This reform focuses on a group of individuals with zero taxable income, because they are very likely to face higher marginal tax rates during the span of the panel. We further invoke the restriction that individuals affected by the reform did not receive benefits in order to exclude long-term benefit recipients. The evidence on mean reversion found in the literature suggests that the incomes of the remaining taxpayers are very upward mobile, indicating the importance of intertemporal tax wedges for these individuals. On average, this reform affects 1,264 tax units in our sample in each of the four years, corresponding to 6,760,126 individuals each year using panel weights.

For this reform we choose the same elasticity scenarios as for Reform 2. Specifically, we consider values for η of 0.1, 0.15, 0.2, 0.3, and 0.4, and for ε_{nm} of 0, -0.01 , -0.02 , -0.05 , and -0.1 . The simulation results are in Table 4.

The reform improves economic efficiency even for very small cross-price elasticities. For instance, the static marginal deadweight loss is estimated to be 0.06% of tax revenue if $\eta = 0.2$. If further $\varepsilon_{nm} = -0.05$, the estimate changes to -0.04% , which is lower by 155%. This illustrates that efficiency is improved when intertemporal substitution is taken into account, even though we are considering a marginal tax increase. It is worth noting that the static marginal deadweight losses in the first row are all positive. Hence, the efficiency

Cross-price elasticity ε_{nm}		Static elasticity η				
		0.1	0.15	0.2	0.3	0.4
0	DWL:	0.04	0.06	0.08	0.11	0.15
	%-dev.:	0	0	0	0	0
-0.01		0.01	0.03	0.05	0.09	0.13
		-61.88	-41.25	-30.94	-20.63	-15.47
-0.02		-0.01	0.01	0.03	0.07	0.11
		-123.76	-82.51	-61.88	-41.25	-30.94
-0.05		-0.08	-0.06	-0.04	0.00	0.03
		-309.39	-206.26	-154.70	-103.13	-77.35
-0.1		-0.20	-0.18	-0.16	-0.12	-0.08
		-618.79	-412.53	-309.39	-206.26	-154.70

Table 4: Marginal phase-out of exemptions (Reform 3).

improvement is not due to a reduction of a static tax wedge (e.g., raising a negative marginal tax), but occurs precisely because of intertemporal income shifting. However, because the reform affects such a small number of individuals and the tax bases are so small, this result is of higher qualitative than quantitative importance. Nevertheless, it highlights that accounting for intertemporal substitution in taxable income can reverse policy recommendations.

Robustness of results

One concern with the above analysis is the assumption of constant intertemporal cross-price elasticities. The evidence on intertemporal substitution deals only with the very short run, typically neighboring years, such that there is little knowledge of cross-price

	Reform 1			Reform 2			Reform 3		
	η			η			η		
ϵ_{nm}	0.3	0.4	ϵ_{nm}	0.15	0.2	ϵ_{nm}	0.15	0.2	
<i>Two cross responses</i>									
DWL:	-0.1	6.07	7.85	-0.02	5.09	6.84	-0.02	0.05	0.06
%-dev.:		13.27	9.95		-3.53	-2.65		-21.30	-15.98
	-0.2	6.78	8.56	-0.05	4.34	6.10	-0.05	-0.00	0.02
		26.53	19.90		-17.65	-13.24		-106.50	-79.88
	-0.4	8.20	9.99	-0.1	3.41	5.17	-0.1	-0.06	-0.05
		53.06	39.80		-35.30	-26.47		-213.01	-159.75
<i>One cross response</i>									
	-0.1	5.73	7.52	-0.02	5.17	6.93	-0.02	0.05	0.07
		7.03	5.28		-1.86	-1.39		-10.84	-8.13
	-0.2	6.11	7.90	-0.05	4.78	6.54	-0.05	0.03	0.05
		14.07	10.55		-9.30	-6.97		-54.18	-40.64
	-0.4	6.87	8.65	-0.1	4.29	6.05	-0.1	-0.00	0.01
		28.14	21.10		-18.59	-13.95		-108.37	-81.28

Table 5: Sensitivity of results to different assumptions about intertemporal responses.

responses beyond a one-year horizon. However, it seems likely that there is considerably less intertemporal substitution between years that are farther apart. We address this concern by running our simulations under different assumptions about the cross-price elasticities. In one scenario, we consider only cross-price responses in two adjoining years. In another scenario, we consider only cross-price responses in one adjoining year. The results are reported in Table 5 for selected elasticities.

Results for two cross responses are reported in the upper half of the table. Similarly, the lower half of the table reports results for one cross response only. In the former

case, we set the cross-price elasticities to zero in all but two contiguous years. Hence, for marginal tax changes in 1997 and 1998 there are cross-price responses in neighboring years, whereas for 1996 and 1999 there are responses in the two following and preceding years, respectively. The results are not much affected by the exact choice of years; only the number of years with cross-price responses matters. Comparing the results to Tables 2–4 shows a dampening of the bias from ignoring intertemporal substitution. While this was to be expected, the biases remain substantial.

A second concern is that benefit phase-out has a large impact on marginal tax rates for low income taxpayers in our sample (see Figure 2 in Appendix C), and that this might influence the results. In results not reported, we have addressed this concern by repeating our simulations taking account of phase-out. This affects the level of the marginal deadweight loss but has no appreciable influence on the bias from ignoring intertemporal substitution.

Our analysis has not explicitly dealt with the role of uncertainty. Uncertainty about future income prospects and tax rates affects both the behavior and welfare of individuals. Most likely, individuals will be less inclined to respond to future tax changes when the changes are uncertain, which would be reflected in lower elasticities of intertemporal substitution. This is an empirical question that only concerns behavior and should be kept in mind when considering the appropriate values of the intertemporal elasticity of substitution. Importantly, the available estimates of the elasticity of substitution reflect the underlying uncertainty in the empirical setting and therefore address this concern

directly. On the other hand, we do not take the welfare consequences of uncertainty into account. This requires assumptions about risk aversion and is beyond the scope of the paper.

To simplify the exposition, our analysis assumes that the interest rate is zero. The analysis can be extended to incorporate discounting but this has little significance for the results. A positive interest rate would affect the level of the estimated marginal deadweight losses, but has no appreciable, nor economically meaningful, effect on the bias from ignoring intertemporal tax wedges.

The main conclusion that emerges from the simulations is the importance of taking into account the effects of a tax reform on intertemporal tax wedges. Even cross-price responses that are, by any standards, small often substantially changes the estimated deadweight loss of raising a marginal tax rate slightly. This insight suggests that the efficiency of the current U.S. tax code may have to be reassessed and, specifically, that the deadweight loss of tax progression is higher than previously thought.

5 Conclusion

It is important to acknowledge that the practice of taxing income at the annual level holds major implications for the study of income taxation. We have shown that a static framework can produce substantially biased estimates of the marginal deadweight loss of income taxation because it ignores intertemporal tax wedges and intertemporal substitution in taxable income. The failure to recognize intertemporal substitution leads to

an underestimate of the marginal deadweight loss when intertemporal tax wedges are increased and an overestimate when intertemporal tax wedges are decreased. In fact, accounting for intertemporal substitution can reverse conclusions about the efficiency of tax reform. A general insight is that the efficiency costs of tax progression are higher once intertemporal responses are taken into account. However, there is still considerable uncertainty about intertemporal responses to tax changes and it is crucial to improve our working knowledge in this area in order to make precise policy recommendations.

An important feature of actual tax systems that has been ignored in the analysis is the taxation of capital income. The literature on income taxation has traditionally focused either on earned income or on capital income, often allowing limited interplay between the two instruments. Indeed, most studies of earned income taxation have relied on static models rendering analysis of capital taxation futile from the outset. Intertemporal tax wedges hold implications for capital income taxation because changes to the timing of income are accompanied by either a savings response or a response in the timing of consumption. Changes in the incentive to save, for example through capital income taxation, can increase the cost of intertemporal substitution of income and help to alleviate the distortions from intertemporal tax wedges. This complex interdependency between the tax incentives for intertemporal income shifting and the efficiency of capital income taxation is an interesting area for future research.

Appendix

A Derivation of the Marginal Deadweight Loss

This appendix contains the derivation of the marginal deadweight formula in Section 2.

We consider a marginal change in the tax rate t^j in a single tax bracket j . From the text, we know that the marginal deadweight loss is expressed by

$$\frac{dDWL}{dt^j} = \sum_{n=1}^N \frac{\tau_n}{1-t^j} z_n \sum_{m \in \Omega} \varepsilon_{nm}, \quad (\text{A-1})$$

where $\varepsilon_{nm} \equiv [(1-\tau_m)/z_n] \partial z_n / \partial (1-\tau_m)$ is the compensated elasticity of taxable income in period n with respect to the marginal net-of-tax rate in period m . The static elasticity is the compensated taxable income response to a permanent change in the marginal net-of-tax rate for a taxpayer who faces the same marginal tax rate, τ , in all periods

$$\eta_n \equiv \frac{1-\tau}{z_n} \frac{dz_n}{d(1-\tau)} = \sum_{m=1}^N \varepsilon_{nm}. \quad (\text{A-2})$$

Using (A-2) in (A-1) gives

$$\begin{aligned} \frac{dDWL}{dt^j} &= \sum_{n=1}^N \frac{\tau_n}{1-t^j} z_n \sum_{m \in \Omega} \varepsilon_{nm} \\ &\quad + \sum_{m \in \Omega} \frac{t^j}{1-t^j} z_m \eta_m \\ &\quad - \sum_{m \in \Omega} \frac{t^j}{1-t^j} z_m \sum_{n=1}^N \varepsilon_{mn} \\ &= \frac{t^j}{1-t^j} \sum_{m \in \Omega} z_m \eta_m \\ &\quad + \sum_{m \in \Omega} \sum_{n \neq m} \left(\frac{\tau_n}{1-t^j} z_n \varepsilon_{nm} - \frac{t^j}{1-t^j} z_m \varepsilon_{mn} \right), \end{aligned}$$

which is (4) from the text.

B General Tax Reform

This appendix considers the marginal deadweight loss from a general change, $d\tau$, in the tax function $T(z)$. Each taxpayer solves the expenditure minimization problem in (1). As in Section 2, τ_n is the taxpayer's marginal tax rate in period n , but now Ω is the set of periods where the marginal tax rate changes following the reform. For each taxpayer we derive the marginal deadweight loss of the reform as in Section 2

$$\begin{aligned} \frac{dDWL}{d\tau} &= \frac{de(T, \bar{u})}{d\tau} - \sum_{n=1}^N \frac{dT(z_n)}{d\tau} \\ &= \sum_{n=1}^N \left\{ \frac{\partial T(z_n)}{\partial \tau} - \sum_{m \in \Omega} \left[\frac{\partial c_n}{\partial (1 - \tau_m)} - (1 - \tau_n) \frac{\partial z_n}{\partial (1 - \tau_m)} \right] \right\} \\ &\quad - \sum_{n=1}^N \left[\frac{\partial T(z_n)}{\partial \tau} - \tau_n \sum_{m \in \Omega} \frac{\partial z_n}{\partial (1 - \tau_m)} \right]. \end{aligned}$$

There are no first order effects on the minimized expenditure from behavioral responses, such that we can write the marginal deadweight loss as

$$\frac{dDWL}{d\tau} = \sum_{n=1}^N \sum_{m \in \Omega} \frac{\tau_n}{1 - \tau_m} z_n \varepsilon_{nm}, \quad (\text{B-1})$$

where $\varepsilon_{nm} \equiv [(1 - \tau_m)/z_n] \partial z_n / \partial (1 - \tau_m)$ is the compensated elasticity of income in period n with respect to the marginal net-of-tax rate in period m .

By inserting the definition of the static elasticity η

$$\eta_n \equiv \frac{1 - \tau}{z_n} \frac{dz_n}{d(1 - \tau)} = \sum_{m=1}^N \varepsilon_{nm},$$

we can rewrite (B-1) as

$$\begin{aligned}
 \frac{dDWL}{d\tau} &= \sum_{n=1}^N \sum_{m \in \Omega} \frac{\tau_n}{1 - \tau_m} z_n \varepsilon_{nm} \\
 &\quad + \sum_{m \in \Omega} \frac{\tau_m}{1 - \tau_m} z_m \eta_m \\
 &\quad - \sum_{m \in \Omega} \frac{\tau_m}{1 - \tau_m} z_m \sum_{n=1}^N \varepsilon_{mn} \\
 &= \sum_{m \in \Omega} \frac{\tau_m}{1 - \tau_m} z_m \eta_m \\
 &\quad + \sum_{m \in \Omega} \sum_{n \neq m} \left(\frac{\tau_n}{1 - \tau_m} z_n \varepsilon_{nm} - \frac{\tau_m}{1 - \tau_m} z_m \varepsilon_{mn} \right), \tag{B-2}
 \end{aligned}$$

which forms the basis for our numerical simulations in Section 4.

C Benefit Phase-Out

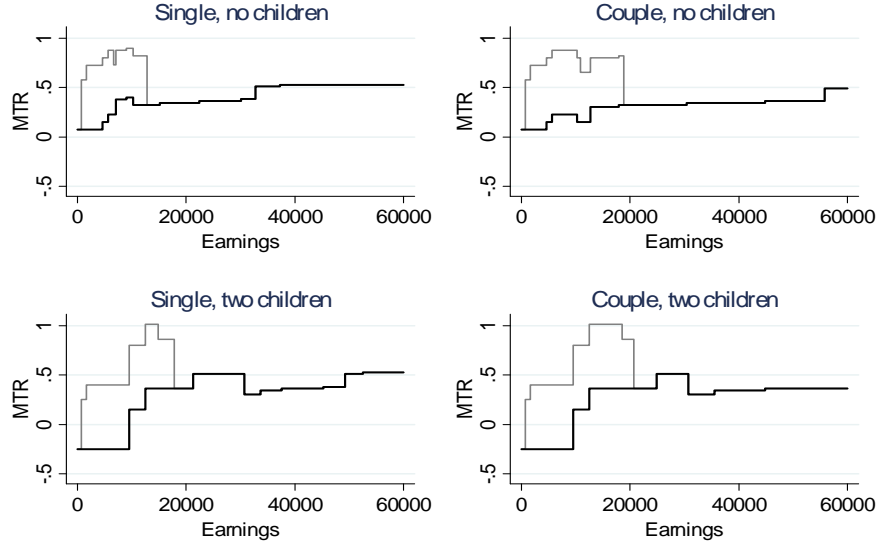


Figure 2: Effective low income marginal tax rates in California, 1999.

The black line displays combined state and federal effective marginal income tax rates. The gray line includes benefit phase-out rates. Source: NBER's TAXSIM model.

References

- [1] Auerbach, A. (1985). “The Theory of Excess Burden and Optimal Taxation.” In Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, Vol. 1, Amsterdam: North-Holland, pp. 61–127.
- [2] Browning, E.K. (1987). “On the Marginal Welfare Cost of Taxation.” *American Economic Review*, 77, pp. 11–23.
- [3] Feldstein, M. (1995). “The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act.” *Journal of Political Economy*, 103, pp. 551–572.
- [4] Feldstein, M. (1999). “Tax Avoidance and the Deadweight Loss of the Income Tax.” *The Review of Economics and Statistics*, 81, pp. 674–680.
- [5] Gaube, T. (2007). “Optimum Taxation of Each Year’s Income.” *Journal of Public Economic Theory*, 9, pp. 127–150.
- [6] Goolsbee, A. (2000). “What Happens When You Tax the Rich? Evidence from Executive Compensation.” *Journal of Political Economy*, 108, pp. 352–378.
- [7] Gruber, J. and E. Saez (2002). “The Elasticity of Taxable Income: Evidence and Implications.” *Journal of Public Economics*, 84, pp. 1–32.
- [8] Harberger, A.C. (1964). “Taxation, Resource Allocation, and Welfare.” In John

- F.Due, ed., *The Role of Direct and Indirect Taxes in the Federal Revenue System*, Princeton: Princeton University Press, pp. 25–70.
- [9] Kopczuk, W. (2005). “Tax Bases, Tax Rates and the Elasticity of Reported Income.” *Journal of Public Economics*, 89, pp. 2093–2119.
- [10] Ljungqvist, M. and K. Ragan (2006). “Labor Supply and the Tax Reform of the Century.” Mimeo.
- [11] Looney, A. and M. Singhal (2006). “The Effect of Anticipated Tax Changes on Intertemporal Labor Supply and the Realization of Taxable Income.” Mimeo.
- [12] MaCurdy, T.E. (1981). “An Empirical Model of Labor Supply in a Life-Cycle Setting.” *Journal of Political Economy*, 89, pp. 1059–1085.
- [13] Saez, E. (2002). “Do Taxpayers Bunch at Kink Points?” NBER Working Paper No. 7366.
- [14] Werning, I. (2007). “Optimal Fiscal Policy with Redistribution.” *Quarterly Journal of Economics*, 122, pp. 925–968.

Chapter 2

Public Provision of Private Goods: Optimal Provision Schemes*

Nicolaj Verdelin
University of Copenhagen and EPRU

August 2008

Abstract

This paper analyzes the optimal design of policy schemes for public provision of private goods under asymmetric information by asking whether provision should be universal or targeted through means-testing. Further, the paper determines the optimal tax treatment of purchases of market alternatives to publicly provided private goods when universal provision is implemented. Two important conclusions emerge from the analysis. First, universal provision can never be justified as an efficient means of redistribution and may diminish the government's scope for redistribution. Instead, the optimum is always *de facto* means-tested. Second, it is shown that the optimal tax policy under universalism involves lump sum subsidies that decrease with income. Top incomes are likely not to be eligible for a subsidy. A tax expenditure type policy that makes private purchases deductible from the income tax base does not implement the optimum. The subsidy result follows from an equivalence between equilibria with means-testing and equilibria with universal provision and subsidization.

Keywords: Universalism; Private goods; Means-testing; Subsidization; In kind transfers
JEL classification: H21; H42

*I am indebted to Jes Winther Hansen and Claus Thustrup Kreiner for very inspirational conversations and comments. All remaining errors are mine. The project has been supported by a grant from the Economic Policy Research Network (EPRN).

1 Introduction

There are numerous examples of governments providing private goods to their citizens. Most prominent are the cases of health care, education and various forms of social insurance but the list is certainly longer, especially in Europe. While some government involvement is common in most developed countries, the range and scope of the government's role differs vastly across countries. In the U.S., the right to public provision tends to be means-tested and limited to fairly small groups, although public schooling is a notable exception. In many European welfare states, on the other hand, the government is typically the main provider of a number of private goods and provision is more often than not available to all citizens.

From the perspective of optimal policy design universal public provision of private goods gives rise to two basic questions. One concerns the justification for government intervention in goods provision and has been studied at least since Musgrave (1959). The other asks whether or when provision should be universally available and has received much less attention. The focus of the present paper is to examine when universalism is part of the optimal policy design.

There may be instances when universal public provision is justified *per se* because private provision is hampered by a global market failure, or because the government has direct preferences over the consumption patterns of all individuals. However, such instances are not the topic of this paper. Instead, the present analysis assumes, as is standard in public finance, that the design of government policies is guided by a desire

to help the least well-off. Thus, government intervention in private goods provision is justified only as a means to help the poor, either directly or through increased redistribution. When the government cannot observe true needs directly (e.g., innate ability), it is forced to rely on an imperfect signal such as income. This raises a basic question of policy design; whether publicly provided goods aimed at the poor should be targeted through means-testing or made universally available.

An argument in favor of universal provision must imply that it gives rise to an equilibrium that is superior to any equilibrium attainable without universal provision. This does not necessarily require that all agents consume the publicly provided good, but the equilibrium must hinge on the good being available to all agents. To be precise, this paper speaks of public provision when the government makes a certain private good available in a predetermined quantity or quality to a possibly limited set of individuals. Often there will be no direct link between individual consumption and individual expenditure. Instead, the program is financed through the tax system and contributions are compulsory. Public provision is *universal* if the good is made available to all members of society. This definition does not invoke any direct restrictions on private provision. A market alternative will be available if it is profitable, i.e., if there is private demand at profitable prices. But universal provision implicitly forces everyone to purchase the publicly provided good.

One of the main criticisms of public provision of private goods is that it violates the principle of consumer sovereignty. It is generally best to leave the choice of consumption

bundle to the individual consumer since she is presumably better informed about her own preferences than the government. With heterogeneous consumers, a policy that forces everyone to buy a certain good in a fixed quality is bound to introduce inefficiencies in provision. The traditional argument in favor of government interference in goods provision refers to public goods and has no bearing on private goods. The private provision of public goods is hampered by free-riding, which leads to inefficiently low provision in the laissez-faire equilibrium. This can be solved either by tax incentives or direct public provision (see, e.g., Diamond, 2006). However, unlike public goods, private goods are both rival and excludable and therefore do not suffer from free-riding.

The fact that government provision comes at a cost does not rule out the scope for beneficial intervention. One such instance is when provision is inefficient in the private equilibrium. This can occur, for instance, when market interaction suffers from asymmetric information. Indeed, it is well established that informational asymmetries may cause markets to collapse entirely. While this *may* justify government provision, there is no general presumption that governments can overcome these inefficiencies if public officials suffer from the same informational barriers as do economic agents.

Paternalism has also been invoked in favor of government provision. If some agents are boundedly rational and tend to, e.g., underconsume certain goods when left to themselves, governments may beneficially intervene by providing the goods in question directly. For instance, social security is often justified by a concern that many may be saving too little for retirement. But when consumers are heterogeneous, with respect to earnings

potential, preferences or behavioral lapses, the benefits of this type of policy typically must be weighted against the cost of forcing people to enter the program.

Finally, in kind transfers have been suggested as an efficient means of redistribution when preferences differ in a systematic way between agents of different ability (Nichols and Zeckhauser, 1982, and Blackorby and Donaldson, 1988). Specifically, the government can achieve more redistribution by including in the transfer package goods that are valued less by the more able, than it can relying solely on a non-linear income tax. It is crucial for this line of reasoning that transfers be means-tested.

Many publicly provided private goods such as education or health care are consumed in discrete quantities. In these cases, the relevant decision margin is usually the quality of the good consumed rather than the quantity purchased (Besley and Coate, 1991). For instance, patients in need of medical surgery value the quality of the procedure, not the number of operations. Similarly, parents do not send their children to more than one school but instead value the quality of teaching. When consumption is discrete, a person who chooses to rely on public provision will not demand the same good from the private market. The formal analysis in this paper treats the case of discrete consumption goods, although the results are also relevant when public provision can be supplemented by private purchases, as long as the publicly provided good and the preferred market alternative are not perfect substitutes.

The present paper shows that universal provision is generally Pareto dominated by means-tested programs. Indeed, the redistributive optimum with public provision of pri-

vate goods is always *de facto* means-tested. Crucially, this implies that the redistributive optimum is unaffected by the preferences of the high-productivity individuals in contrast to the results in Blomquist and Christiansen (1995) who restrict the instruments available to the government. It is possible to implement the optimum under universal provision with the use of subsidies on private purchases of market alternatives to the public quality. But this essentially amounts to means-testing as individuals above a certain income level are induced to seek private provision. Indeed, there is a formal equivalence, in terms of equilibrium outcomes, between means-tested equilibria and equilibria with universal provision and subsidy-induced separation of types, which implies that the above result makes the strongest possible case against universal provision. The analysis highlights that the common concern that means-testing provokes undesirable earnings distortions, because an earnings test functions just like a tax, is not relevant when transfers are in kind. The only assumption needed for this conclusion is that the government be able to compensate, through the income tax, those excluded from a program.

The results in this paper have important implications for the design of the modern welfare state but also for optimal tax policy in societies that have implemented universal public provision of private goods. Indeed, regardless of society's reasons for publicly providing private goods the spread of such policies raises the important question of the appropriate tax treatment of private purchases of substitute goods. Such purchases are widely observed. For instance, parents often send their children to private schools even though a free public school is available. Similarly, private hospitals play an increasing role

in countries that have universal public health care. This paper shows that it is optimal to subsidize purchases of market alternatives to publicly provided private goods. The optimal subsidy decreases with income when the good is normal. However, the optimum may be implemented with a uniform subsidy as long as the income tax is sufficiently flexible. A policy that makes private expenditures deductible from taxable income does not implement the optimum.

There is no formal modelling of the justification for public provision to low incomes in this paper, but I view this feature as a strength, not a weakness. The choice between universal and means-tested provision should not be affected by the specific reason for government intervention in the first place, unless public provision is justified by a global market failure that affects all citizens alike, or the government has direct preferences over individual consumption patterns. Thus, the analysis remains valid for *any* justification for public provision that does not require the participation of the entire population a priori.

There is a small existing literature dealing with public provision of private goods. Besley and Coate (1991) argue that a number of publicly provided goods are consumed in discrete quantities. They show that universal provision can redistribute from rich to poor when the level of public provision is set low enough that the rich opt for private provision. Blomquist and Christiansen (1995) and Boadway and Marchand (1995) analyze optimal public provision when an optimal non-linear income tax is implemented. Blomquist and Christiansen (1998) consider the design of public provision schemes when the government

can disallow benefit recipients from supplementing the public provision. Besley (1991) argues that it is always optimal to have at least small user charges on the consumption of publicly provided private goods.

The paper is organized as follows. Section 2 sets up a model with two types and examines the case for universal provision and the optimal tax treatment of purchases of publicly provided goods. Section 3 discusses the generality of the results and concludes. A general model with continuous ability types is presented in the Appendix.

2 Universal Public Provision of a Discrete Good

This section presents a formal analysis of public provision of a discrete consumption good with variable quality. When there is public provision of a discrete good, people are essentially forced to pay the public provider whether or not they consume the good in question. If some consumers are not satisfied with the quality of the government provision, the only alternative is to switch entirely to a private provider. This amounts to giving up on a consumption good that is free at the margin. In this sense public provision of a private good vastly distorts the relative price of the public quality compared to other quality levels.¹ At the same time, it reduces the disposable income of all citizens forced to enter the program.

Following Besley and Coate (1991) consumers allocate their spending across a numeraire good c and the quality q of a single unit of the good x . All consumers are

¹The publicly provided good need not be free. What is required for the analysis is that the good be subsidized and that the subsidy is paid for by everyone forced to enter the program.

assumed to have smooth, well-behaved preferences for the numeraire and the quality of good x given by $u(c, q)$. As in Stiglitz (1982), the economy consists of two types of agents, differing only by their earnings ability. Low types, denoted L , have ability n_L and high types, denoted H , have ability n_H . An agent with ability n must exert a level of effort corresponding to y/n to generate income y . The share of low types is given by γ . The preferences of agent $i = L, H$ are given by

$$U(c, q, y) = u(c, q) - v(y/n_i),$$

where $u_c, u_q > 0$, $u_{cc}, u_{qq} < 0$, $u_{cq} > 0$, and $v', v'' > 0$. These assumptions ensure that both c and q are normal goods, which implies that the optimal choice of quality is higher for type H than for type L . I further assume that conditions are met to always ensure an interior solution.

It is the government's objective to supply good x to people of *type L*. The underlying justification for the government's intervention in goods provision is not explicitly modelled but a few possible reasons were mentioned in the Introduction. Thus, the analysis builds on the (weak) assumption that the choice between universal provision and means-testing is not affected by the reason for supplying good x in the first place. Since idiosyncratic productivities are unobservable to government officials, public provision cannot be made directly dependent on the individual's type. Instead, public services must be either means-tested or available to all citizens as is the case with universal provision.

Under universal provision, the government provides good x at a fixed quality level q_g implying that all agents can consume the quality level q_g at zero (private) marginal cost.

The government has access to the same technology as the private sector, implying that the cost per citizen of supplying q_g is πq_g , where π is the relative price of quality of good x . However, the public provider is only an imperfect alternative to the private market since agents who choose to rely on public provision must consume the quality level q_g irrespective of their own preferred choice of quality. As a result, the entire consumption bundle of agent i is affected by the choice of provider of good x . In the following the subscript ‘ g ’ denotes choices when agent i relies on public provision and the subscript ‘ p ’ refers to the choices of agent i when she seeks a private provider.

The government operates a (possibly) non-linear income tax $T(y)$ that finances all public expenditures and potentially redistributes income across the two types. The tax function need not be the result of welfare maximization. The budget constraint of agent i is

$$z_i \equiv y_i - T(y_i) = \begin{cases} c_i & \text{for } q_i = q_g \\ c_i + \pi q_i & \text{for } q_i \neq q_g \end{cases},$$

where z_i is the after-tax income of agent i . The budget set of consumer i depends on whether she settles for public provision at the quality level q_g , or opts for a private alternative. As a reflection of the discrete nature of good x , the consumer bears the entire cost (equal to πq_i) of deviating from the public provision level.

Under universal provision, agent i chooses private provision if and only if she values the freedom to choose her preferred quality of good x highly enough to outweigh the cost

of replacing the publicly provided level²

$$u(z_{ip} - \pi q_{ip}, q_{ip}) - v(y_{ip}/n_i) > u(z_{ig}, q_g) - v(y_{ig}/n_i),$$

where q_{ip} is chosen to satisfy $u_q(c_{ip}, q_{ip}) = \pi u_c(c_{ip}, q_{ip})$ and $y_{ij}, j = g, p$, satisfies $v'(y_{ij}/n_i) = (1 - \tau_{ij}) n_i u_c(c_{ij}, q_{ij})$ with τ_{ij} the marginal tax rate for type i when choosing provider j and $q_{ig} = q_g$.³ Naturally, no-one opts out of the public program if the public quality level exceeds their own preferred quality. Lemma 1 shows that high ability types are always more likely to opt for private provision (all proofs are in Appendix A).

Lemma 1. *In any equilibrium with universal public provision, type H is always more likely to choose private provision than type L when q is a normal good. Specifically, in any equilibrium where type H is just indifferent between public and private provision, type L settles for the public provider. Further, for any non-degenerate tax function $T(\cdot)$, the earnings of type H are higher whenever she chooses private provision.*

When setting the level of public quality, the government initially has the interests of type L in mind. However, when the provision scheme is universal, the public supply may prevent a private market from arising, if the quality is sufficiently high, i.e.,

$$u(z_{Hp} - \pi q_{Hp}, q_{Hp}) - v(y_{Hp}/n_H) < u(z_{Hg}, q_g) - v(y_{Hg}/n_H). \quad (1)$$

If this condition is met even when the public quality is chosen to serve the interests of

²This expression ignores the budget effect from the savings in public expenditures on x because each agent is assumed atomistic. Further, it is assumed that, when an agent is indifferent between public and private provision, she opts for public provision. This assumption is completely innocuous but it serves to make the proof of Proposition A1 (in Appendix) for the continuous type model easier.

³These conditions apply when the tax function is differentiable.

type L only, the government faces the choice between either (a) forcing people of high ability to accept the public quality (possibly adjusting q_g upwards to better service the needs of the expanded user group) or (b) settle for a lower public quality that does not induce type H to opt for the public program. In the former case, while the high types settle for public provision, the quality provided is inevitably too low from their viewpoint.

That is

$$u(z_{Hg} - \pi(q - q_g), q) \text{ is increasing in } q \text{ at } q = q_g \quad (2)$$

This implies that when public provision discourages high ability individuals from seeking private providers, their quality choice is distorted (as will be the quality choice of type L if provision is adjusted to partly reflect the wishes of both types). The two conditions (1) and (2) if met when evaluated at the value of q_g that best serves the interest of type L act as constraints on the government problem. When this is the case the government cannot implement universal provision at the preferred level of q_g without distorting the consumption choice of type H . In this case, I will speak of *universal provision with public sector dominance*. This concept refers to an equilibrium occurrence that acts as a constraint on the government optimum, not to the optimum itself. From now on, I am interested in the optimal provision scheme when the problem of public sector dominance is binding. The section thus proceeds by examining whether the two options (a) and (b) mentioned above truly represent the best options available to the government. The finding is that in fact the government can do better by introducing means-tested provision.

The distortion of the consumption decision represents the central cost associated with

universal public provision of private goods consumed from a single provider. It arises because of the discrete nature of good x , which implies that the public quality level cannot be supplemented through market purchases. Therefore, the public quality must strike a balance between the diverging demands of different individuals. When the public supply is financed through taxation and made available at no or very low marginal cost, consumers have a strong incentive to choose public provision even if the public quality level does not match their preferred choice.

An alternative to universal provision applies a means-test to determine eligibility for the public supply of x . In this case, q_g is only available to individuals with income below a certain threshold \bar{y} . When \bar{y} is set below the equilibrium income of type H but above that of type L there is separation of types in the demand for x . I will assume from now on, when referring to equilibria with means-testing that there is separation of types in the demand for x such that only type L seeks the public provider. Let $T_{mt}(\cdot)$ denote the tax function in a means-tested equilibrium, and $T_u(\cdot)$ the equilibrium tax function with universal provision. Further, let \tilde{y} denote the income of the highest-earning individual that seeks public provision in the universal equilibrium. Finally, let $1(y > \bar{y})$ be an indicator function that takes on the value 1 when y exceeds the income limit for the means test, \bar{y} , and 0 otherwise, and let $1(y > \tilde{y})$ be similarly defined for the threshold \tilde{y} .

The self-selection properties embodied in Lemma 1 imply that the government can induce *effective targeting* of the public provision by subsidizing, at an appropriate rate, purchases of market alternatives to the public q . This further implies that any equi-

librium with means-testing can be imitated by an equilibrium with universal provision and subsidies to private purchases of x . Define, for any q_g the *equivalent subsidy* $s(q_g)$ satisfying

$$u(z_{Hp} - \pi q_{Hp} + s(q_g), q_{Hp}) - v(y_{Hp}/n_H) = u(z_{Hg}, q_g) - v(y_{Hg}/n_H) + \varepsilon, \quad (3)$$

for some $\varepsilon > 0$ and where the left-hand side is evaluated at the means-tested equilibrium allocations (which induce separation in the demand for x) that we are trying to imitate and the right-hand side is evaluated at the initial equilibrium with universal provision and no subsidization. When the equivalent subsidy is implemented under universal provision (possibly including a change to the income tax function), the same allocations as in the means-tested equilibrium obtain. Thus, we can imitate the means-tested equilibrium without introducing a formal means test. Correspondingly, any equilibrium with universal provision that induces separation in the demand for x , through subsidization at a rate s , can be implemented as an equilibrium with means-testing (without subsidization). In this case, the subsidy in place under universalism constitutes the necessary compensation to those that fail the income test under means-testing.⁴

Proposition 1. *Any equilibrium with means-testing can be replicated by an equilibrium with universal provision that implements the equivalent subsidy. The imitating universal equilibrium is supported by a tax function satisfying the relationship $T_u(y) = T_{mt}(y) + s(q_g) \cdot 1(y > \bar{y})$ for any y . Similarly, any equilibrium with universal provision that induces*

⁴Naturally, a means-tested imitating equilibrium can also be implemented by maintaining the subsidy in place in the universal equilibrium. In fact, since the universal equilibrium already induces separation, a formal means-test can be implemented without changing the equilibrium allocations.

separation of types can be replicated by a means-tested equilibrium that eliminates the subsidy s (if in place), implements a means-test at the income level \tilde{y} , and is supported by $T_{mt}(y) = T_u(y) - s \cdot 1(y > \tilde{y})$ for any y .

Whenever an equilibrium induces effective targeting of the public supply, I will speak of *de facto means-testing*.⁵

Proposition 1 clarifies that a justification for universalism amounts to arguing that subsidy-induced separation of types is undesirable. This would only be the case if public sector dominance of the market for x enhances the redistributive scope of the government. According to Proposition 1, the traditional concern that means-testing provokes undesirable earnings distortions is not warranted as long as the tax system is sufficiently flexible to compensate those not included in the public program. Even so, the required compensation transfers resources to the rich, which may be undesirable from a redistributive point of view. Thus, the question remains if the government would rather solve the problem of public sector dominance by adjusting the level of q_g downwards or simply by including type H in the public program. Proposition 2 shows that both of these alternatives are dominated by *de facto means-testing*.

Proposition 2. *The social optimum implements de facto means-testing and sets q_g to service the interests of type L only.*

⁵It might seem that formal means-testing is more cost-efficient than universalism because public provision becomes less attractive when there is an eligibility cap on earnings. While this is true it does not imply that means-testing can attain a better equilibrium than universalism with subsidization. Indeed, the scope for redistribution is determined by the utility of type H when she mimicks the low type, which is independent of the means test.

Proposition 2 demonstrates that the social optimum is unaffected by the problem of public sector dominance. Governments that implement public provision of a private good with the objective of aiding, directly or through increased redistribution, those less well off should set the level of public provision according to the interests of the target group only, even if provision is made universal. It may be optimal to adjust the quality level if there are innate differences in preferences between types as shown by Nichols and Zeckhauser (1982) and Blackorby and Donaldson (1988), but only when the adjustment serves the interests of type L by increasing redistribution. The most direct implementation of the social optimum applies a means-test to ensure that the public q is only provided to those individuals whose earnings do not exceed the equilibrium earnings of type L . However, as Proposition 1 shows it is also possible to implement the optimum under universal provision by subsidizing private purchases of market alternatives. Further, the two regimes have the same informational requirements. However, since public provision is effectively targeted at type L in both cases, it is mere semantics to say that provision is universal.

The lesson from Proposition 2 is that universalism may in fact diminish the government's scope for redistribution. The distortion of the quality choice of type H when she is included in the public program lowers her utility and leads her to work less. This loss of utility brings no benefit to type L and thus simply reduces the amount of redistribution the government can achieve. Alternatively, setting q_g lower than would otherwise be optimal hurts type L . Instead, the government should set q_g optimally, target provision

at type L and compensate type H . Such a policy leaves the utility of type H unchanged, raises the well-being of type L and saves government funds because the required compensation is less than the savings to the public program.

The conclusion that the redistributive optimum with universal provision is unaffected by the preferences of type H for public versus private provision, as long as the government can subsidize private purchases, is in contrast to Blomquist and Christiansen (1995). They analyze the general redistributive optimum with universal public provision and conclude that it may be optimal to include type H in the public program or that the redistributive optimum may be affected by the need to ensure that only type L opts for public provision. The reason for these results is that Blomquist and Christiansen do not consider the option of subsidizing private purchases of the publicly provided good nor of making public provision means-tested (which, as we have seen, amount to the same). Instead, they conclude that it may be optimal to introduce a negative marginal tax on type H when she is included in the public program. As the logic of Proposition 2 demonstrates, this result is driven by a desire to correct the earnings distortion introduced by the distortion of the quality choice. But the suggested remedy is only second best compared to the optimal policy of de facto means-testing. Thus, the well-known result that the earnings decision of type H is undistorted in the optimum carries over to cases with public provision of private goods.

The argument against universalism also survives when type H opts for private provision even without compensation. In this case, the same equilibrium can be attained by

limiting access to the publicly provided good to those who actually consume it (type L). This does not affect the IC constraint of the high types who therefore continue to opt for private provision. Indeed,

$$u(c_H, q_H) - v(y_H/n_H) \geq u(c_L, q_g) - v(y_L/n_H)$$

is unchanged whether or not type H has access to the public q provided $q_H \neq q_g$. Besley and Coate (1991) have argued that universal provision can achieve redistribution whenever the public quality is sufficiently low to induce the more able types to voluntarily opt out of the public program. While this is certainly true, the argument above makes it clear that a concern for redistribution cannot in itself justify universalism. Besley and Coate, themselves, are well aware that in kind transfers do not constitute an efficient means of redistribution unless there are innate differences in preferences that separate low from high types. And, even then, the benefits must be targeted at the poor to achieve redistribution, either by law through means-testing or by fact through voluntary separation. In either case, this does not present an argument in favor of universal provision.

In terms of tax policy, Proposition 2 provides guidance to governments that have implemented universal public provision of, e.g., health care or primary education on how to treat private purchases of such goods. Indeed, Corollary 1 shows that, when the problem of public sector dominance applies, the government can achieve a Pareto improvement by subsidizing purchases of market alternatives to the public q at a rate that is less than the cost of servicing type H in the public program.

Corollary 1. *Any equilibrium with universal public provision of good x and public sector dominance can be improved upon by subsidizing purchases of market alternatives at a lump sum rate $s < \pi \cdot q_g$. The subsidy is just sufficient to induce type H to opt for private provision.*

The intuition underlying the result is straightforward: the subsidy required to induce type H to opt for private provision is less than the reduction in public expenditures on q because the reform eliminates a distortion of the quality choice of type H . Thus, the proposed subsidy saves government funds and increases the redistributive scope of the government. Alternatively, a better equilibrium arises if individuals with equilibrium earnings in excess of the equilibrium income of type L are denied access to the public x in return for a decrease in their income tax liability corresponding to the subsidy needed to induce separation in the demand for x . Note that the optimal subsidy does not depend on the expenditures of type H when she chooses private provision. This is because the subsidy only needs to correct a wedge between public and private provision. If a traditional subsidy were implemented such that the value of the subsidy is $s = s' \pi q_p$, the quality choice of agents seeking private provision would be distorted. This is uncalled for since the government's savings when an agent opts for private provision are independent of the individual's private quality choice.

In Appendix B, the results in this section are generalized to a setting with continuous ability types. The benefit of the general approach is that it allows a more complete understanding of the optimal subsidy. Specifically, Proposition A1 in Appendix B demonstrates

that the *effective* subsidy decreases with income. Also, no subsidy should be granted to individuals whose income is high enough that they are willing to opt for private provision even without tax incentives. The government can then always improve upon an equilibrium with universal provision by subsidizing the private purchases of x for all agents not belonging to the target group. Since the preference for quality monotonically increases with income, the subsidy required to induce a shift to private provision is smaller the higher is the individual's income.

There are many ways to implement the optimal policy. The most direct way is to link the size of the lump sum subsidy to income. However, this calls for a possibly complicated formula for calculating the subsidy which is likely to decrease continuously with income. Alternative and more practical implementations would apply either a single lump sum subsidy irrespective of income or a stepwise decrease of the size of the subsidy. This requires that the subsidy at any stage be determined by the preferences of the least able among the eligible agents. Further, less variation in the subsidy requires more adjustment of the income tax, which must be sufficiently flexible to correct for the excess transfer awarded to higher incomes. Yet another possible implementation exploits the simple relationship between income and spendings on q when private provision is chosen. In this case, the subsidy is instead made dependent on spendings on q such that lower spendings tricker higher subsidies (in absolute terms).

Importantly, we know from Proposition 1 that the social optimum can be directly implemented through a formal means test. Proposition A1 then gives the adjustments

to the income tax function that are necessary to compensate those who fail the income test. And these compensations decrease with income because the alternative of private provision is relatively more attractive to people with higher incomes.

The crucial insight from the general model is not the various implementation methods, but rather the conclusion that the subsidy decreases with income. An important corollary is that the optimum is *not* implemented by a tax expenditure type policy in which spendings on private purchases of q , e.g., through one's employer, may be subtracted from taxable income. Such a policy yields higher subsidies to higher incomes and to higher spendings on q .⁶ This also highlights the nature of the subsidy argument. The purpose of subsidizing is not to let the induced public savings accrue fully to individuals opting for private provision, but rather to induce to exit those individuals who may be pushed to seek private providers at a lower cost than the cost of including them in the public program.

Proposition 2 and Corollary 1 are robust to a number of variations to the underlying model:

Consumption technology. The above results carry over to situations when good x is consumed in more than one unit (extending to the case of continuous consumption) and possibly from different suppliers as long as the quality of the service provided is variable. If the government supplies a fixed quality q_g that differs from the optimal choice of type H , then the choices of type H are distorted when she chooses to (at least partly)

⁶This type of policy is currently in place in Denmark.

rely on government provision. The severity of the distortion depends on the number of units consumed from the public supplier and the ability of the consumer to compensate for the insufficient quality of the inframarginal units by purchasing a higher quality at the margin than would otherwise be chosen. Also, note that even if x is consumed in varying quantities, as is the case with health care for instance, and the poor consume more than the rich, universal public provision of x continues to be an inefficient means of redistribution because the quality choice of the rich is distorted.⁷

Preference heterogeneity. The case against universalism survives if good x is inferior or preference differences between types L and H imply that L consumes a higher q than does H . It is easily seen that the relation $s < \pi q_g$ does not depend on the quality choice of type H exceeding q_g . Thus, a policy that abandons universalism, implements the preferred quality of type L , and compensates type H at the rate s , corresponding to the equivalent subsidy, saves government funds. More generally, the results carry over to the general preference specification $U_i = u(c, q, y, n_i)$ as long as conditions are met to ensure that types can still be separated using income-dependent subsidies. In this case, the government may exploit the preference differences between types, including any correlation between preferences for q and work effort to discriminate true type L 's from masquerading type H 's. This increases the scope for redistribution as pointed out by Nichols and Zeckhauser (1982) and Blackorby and Donaldson (1988) but does not make

⁷An interesting question concerns the design of the optimal subsidy when good x is consumed in more than one unit. In this case, the subsidy will depend on the extent of private replacement of public services but the precise relationship is likely to be complicated.

a case for universal provision.

Uncertainty about preferences. When the government is uncertain about the preferences of its citizens, this may affect the distribution of the benefits from moving away from universal provision but does not affect the size of the potential gain. Any compensation that is less than the cost of the public provision represents a Pareto improvement. The public administration may then use whatever information it possesses to identify the lowest amount that compensates the rich for the loss of eligibility for the publicly provided good.

3 Concluding Remarks

The superiority of means-testing over universal provision (with public sector dominance) depends crucially on the ability of the government to compensate consumers through the income tax. Thus, while the income tax need not be optimized initially, it is necessary that the government have sufficient flexibility in adjusting the tax function. Further, this line of reasoning clearly does not translate to cash benefits which, when targeted at certain income groups, by definition introduce earnings distortions. But cash transfers are also simply part of the tax-transfer schedule. However, the (weak) dominance of means-testing does extend to any type of in kind transfer, whether the good is consumed in discrete or continuous quantities and whether the good comes in different qualities. As long as the means test applies the same income measure as the income tax and the income tax can be adjusted to compensate the income groups left out of the provision

scheme, universalism cannot improve the equilibrium.

Social security is an example of a continuous consumption good that is provided by governments in most OECD countries. At the same time, the case of social security highlights the importance for the logic of Corollary 1 of a flexible income tax. It is commonly argued that universal social security has the merit that it avoids distorting individual savings (see, e.g., Feldstein, 2005). This argument is based on the implicit assumption that the income test uses retirement income or asset holdings to determine eligibility. In this case social security is equivalent to an income transfer. However, in principle targeting need not be distortionary if the earnings test applies an income measure that depends on earnings prior to retirement *and* the income tax can be made dependent on the same income measure. Further, to make sense, this income measure must be different from lifetime income. Whether the government would be interested in conditioning social security on such an income measure depends on the government's motivation for providing social security in the first place. A different issue is whether the government is even interested in making social security non-distortionary, or if this provides an efficient means of redistribution (see Diamond, 2003, and Golosov and Tsyvinski, 2006).

Naturally, the superiority of means-testing builds on the assumption that there are no direct reasons to supply the private good to all citizens. The Introduction alluded to the most common justifications for government intervention in the provision of private goods: asymmetric information and market failure, paternalism, and redistribution. The above analysis clarifies that a concern for redistribution can never justify universalism. As to

market failure, there may be reasons for the government to introduce universal public provision of q if a private market for q is not sustainable for any income group . But this requires the existence of a global market failure. Focusing on the case of insurance markets, Rothschild and Stiglitz (1976) have demonstrated that there may be situations under asymmetric information where the laissez-faire equilibrium is Pareto dominated by a privately non-sustainable pooling equilibrium.

As to paternalism, it is clear that if agents from all income groups suffer from a behavioral weakness, which leads them to, e.g., underconsume a certain good, the government may find it optimal to supply the good to all groups. While there may be examples of such general weaknesses, Bertrand et al. (2003) argue that bounded rationality is likely to have more severe consequences for the welfare of poor individuals.

There could be other reasons to opt for universal provision directly related to the attributes of the good in question. It is possible that the government has distributional preferences over certain goods such as health care. This might reflect social norms that value the health of different citizens equally, irrespective of individual preferences, or a discrepancy between willingness-to-accept and willingness-to-pay for health care at the individual level. The latter may lead poor citizens not to purchase high quality health care although they would not re-sell the right to high quality treatment at the going prices if it were given to them. Now suppose the supply of quality is limited, e.g., because not all doctors are equally proficient. Then the government may prefer an allocation where everyone has access to the high quality services in some cases to one where this privilege

is limited to the wealthiest individuals. In this case, however, we should see outright bans on private supplies of health care. And the central lesson from this paper is that social preferences for universalism come at a cost in terms of reduced efficiency.

Finally, there may be political reasons for universal provision. It is sometimes argued that means-testing may limit public support for the welfare state. If so, universalism may be a necessary means, even if the government is only interested in aiding those least well-off.

A Appendix A

Proof of Lemma 1. Let $Q_i \equiv \{q_g \mid u[z_{ip} - \pi q_{ip}, q_{ip}] - v(y_{ip}/n_i) > u[z_{ig}, q_g] - v(y_{ig}/n_i)\}$ be the set of q_g 's for which agent i opts for private provision. Note that $q_{ip} > q_g$ for any $q_g \in Q_i$. Consider some $q_g \in Q_L$. If the equilibrium is such that either type L or type H masquerades with public provision, the result follows immediately. If not, there are two cases to consider:

a) $y_{Lp} \geq y_{Hg}$. It follows from optimization by type L and revealed preference that

$$\begin{aligned} u(z_{Lp} - \pi q_{Lp}, q_{Lp}) - v(y_{Lp}/n_L) &> u(z_{Hg}, q_g) - v(y_{Hg}/n_L) \\ &\Downarrow \\ u(z_{Hp} - \pi q_{Hp}, q_{Hp}) - v(y_{Hp}/n_H) &\geq u(z_{Lp} - \pi q_{Lp}, q_{Lp}) - v(y_{Lp}/n_H) \\ &> u(z_{Hg}, q_g) - v(y_{Hg}/n_H) \end{aligned}$$

since for $y' \geq y$

$$\frac{d[v(y'/n) - v(y/n)]}{dn} = v'(y/n) \frac{y}{n} - v'(y'/n) \frac{y'}{n} \leq 0.$$

b) $y_{Lp} < y_{Hg}$. By revealed preference

$$\begin{aligned} u(z_{Lp} - \pi q_{Lp}, q_{Lp}) - v(y_{Lp}/n_L) &> u(z_{Lg}, q_g) - v(y_{Lg}/n_L) \\ &\geq u(z_{Lp}, q_g) - v(y_{Lp}/n_L). \end{aligned}$$

The complementarity of q and c in $u(\cdot)$ ensures that for $q \geq q_g$

$$\frac{\partial [u(z - \pi q, q) - u(z, q_g)]}{\partial z} = u_c(z - \pi q, q) - u_c(z, q_g) > 0.$$

Since $z_{Hg} > z_{Lp}$ and q is a normal good it must be that

$$\begin{aligned} u(z_{Hp} - \pi q_{Hp}, q_{Hp}) - v(y_{Hp}/n_H) &\geq u(z_{Hg} - \pi q, q) - v(y_{Hg}/n_H) \\ &> u(z_{Hg}, q_g) - v(y_{Hg}/n_H), \end{aligned}$$

where the first inequality is due to revealed preference. The arguments in a) and b) hold for any feasible choices for type H , and in particular also if type H masquerades with private provision. Since utility is continuous in q_g , it follows that $Q_L \subset Q_H$.

Since $q_{Hp} > q_g$ it follows immediately that $u_c(z_{Hg} - \pi q_{Hp}, q_{Hp}) > u_c(z_{Hg}, q_g)$. As a result, $z_{Hp} \geq z_{Hg}$ with strict inequality if there is some $y > y_{Hg}$ and $z > z_{Hg}$ such that $u(z - \pi q_{Hp}, q_{Hp}) - u(z_{Hg}, q_g) > v(y/n_H) - v(y_{Hg}/n_H)$.⁸ \square

Proof of Proposition 1. Consider an equilibrium with means-testing. For earnings below the means test \bar{y} , the tax function in the modified equilibrium with universal provision is $T_u(y_i) = T_{mt}(y_i)$. Therefore the set of available gross income–consumption pairs (y_i, c_i) is the same as in the equilibrium with means-testing. For incomes above \bar{y} the set of available gross income–consumption pairs is the same as in the equilibrium with means-testing when agent i opts for private provision since $c_i = y_i - T_u(y_i) = y_i - T_{mt}(y_i) + s(q_g) - s(q_g) = y_i - T_{mt}(y_i)$. When agent i opts for public provision, consumption at any income level above \bar{y} is less than in the equilibrium with means-testing since $c_i = y_i - T_u(y_i) = y_i - T_{mt}(y_i) - s(q_g)$. Thus, any agent who chose $y_i \leq \bar{y}$ and public provision in the equilibrium with means-testing does so in the modified equilibrium

⁸When the tax function is differentiable, z_{Hp} satisfies $v'(y_{Hp}/n_H) = (1 - \tau_{Hp}) n_H u_c(z_{Hp} - \pi q_{Hp}, q_{Hp})$ which clearly implies $z_{Hp} > z_{Hg}$.

with universal provision. And any agent who chose $y_i > \bar{y}$ and private provision in the equilibrium with means-testing does so in the imitating equilibrium as well. Since type L satisfies the former condition, and type H the latter, the equilibrium outcomes in the modified equilibrium with universal provision are identical to the equilibrium outcomes with means-testing. A similar line of reasoning shows that any equilibrium with universal public provision that induces separation in the demand for x through a subsidy (possibly equal to 0) can be replicated by an equilibrium with means-testing without subsidization.

□

Proof of Proposition 2. Let s^* be a lump sum subsidy on private purchases of q that is just sufficient to induce high types to seek private provision when their income is kept constant. Such a subsidy satisfies

$$u(z_{Hg} - \pi q_p(s^*) + s^*, q_p(s^*)) = u(z_{Hg}, q_g) + \varepsilon,$$

for some $\varepsilon > 0$, and where $q_p(s^*)$ is the optimal choice of quality for type H given income z_{Hg} and the subsidy s^* . It follows from (1) that s^* exists and $s^* > 0$. Further, whenever $u[z_{Hg} - \pi(q - q_g), q]$ is increasing in q at $q = q_g$, it is true that (a) $q_p(s^*) > q_g$ and (b) $s^* < \pi q_g$:

(a) Assume $q_p(s^*) \leq q_g$. It follows that $s^* \geq \pi q_p(s^*)$, but then $q_p(s^*)$ cannot be optimal, since

$$\arg \max_q u\{z_{Hg} - \pi[q - q_p(s^*)], q_p(s^*)\} > q_g$$

by (2). Thus, $q_p(s^*) > q_g$.

(b) Assume $s^* \geq \pi q_g$. From the definition of s^* it follows that $\varepsilon > 0$ exists such that

$$\begin{aligned}
 u(z_{Hg}, q_g) &= u(z_{Hg} - \pi q_p(s^*) + s^*, q_p(s^*)) - \varepsilon \\
 &\geq u(z_{Hg} - \pi q_p(s^*) + \pi q_g, q_p(s^*)) - \varepsilon \\
 &> u(z_{Hg} - \pi q_g + \pi q_g, q_g), \quad \text{as } q_p(s^*) \neq q_g \\
 &= u(z_{Hg}, q_g),
 \end{aligned}$$

which is a contradiction. Thus, $s^* < \pi q_g$.

Let $s = s^*$ and consider a modification of the initial tax function such that the new tax function $T'(y)$ offers only two contracts: $[y_{Lg}, T(y_{Lg})]$ and $[y_{Hg}, T(y_{Hg})]$. Then, by Lemma 1, type L continues to opt for public provision and the contract $[y_{Lg}, T(y_{Lg})]$. Furthermore, type H chooses private provision and the contract $[y_{Hg}, T(y_{Hg})]$ since

$$\begin{aligned}
 &u[z_{Hg} - \pi q_{Hp}(s^*) + s^*, q_{Hp}(s^*)] - v(y_{Hg}/n_H) \\
 &\geq u(z_{Lg}, q_g) - v(y_{Lg}/n_H) \\
 &> u[z_{Lg} - \pi q_{Lp}(s^*) + s^*, q_{Lp}(s^*)] - v(y_{Lg}/n_H),
 \end{aligned}$$

again by Lemma 1. Thus, in the new equilibrium the indirect utilities of both types are the same as in the initial equilibrium.

The change in government revenue from the initial to the new equilibrium is

$$\begin{aligned}
 R' - R &= \gamma T(y_{Lg}) + (1 - \gamma) T(y_{Hg}) - \gamma \pi q_g - (1 - \gamma) s^* \\
 &\quad - [\gamma T(y_{Lg}) + (1 - \gamma) T(y_{Hg}) - \pi q_g] \\
 &= (1 - \gamma) (\pi q_g - s^*) \\
 &> 0
 \end{aligned}$$

from (b) above. This leaves scope for a Pareto improvement. Since a Pareto improvement materializes from any equilibrium with public sector dominance, the social optimum must involve separation in the demand for x . It then follows from Proposition 1 that the social optimum entails de facto means-testing. Further, when only type L chooses public provision, q_g should be set to serve the interests of type L only. \square

Proof of Corollary 1. Follows immediately from the first part of the proof of Proposition 2. \square

B Appendix B

This appendix extends the results of Section 2 to continuous productivity types. Every agent i is identified by her productivity n_i which is drawn from the set $[0; \infty]$. The density of agents with productivity n is $f(n)$ and the number of persons with productivity no higher than n is $F(n)$. I assume that there is no bunching. The level of public provision is q_g and the marginal user of the public service is $\bar{n}(q_g)$ satisfying

$$u(z_{\bar{n}p} - \pi q_{\bar{n}p}, q_{\bar{n}p}) - v(y_{\bar{n}p}/\bar{n}) = u(z_{\bar{n}g}, q_g) - v(y_{\bar{n}g}/\bar{n}). \quad (\text{B-1})$$

It follows that $q_{\bar{n}p} > q_g$ and thus

$$u(z_{\bar{n}g} - \pi(q - q_g), q) \tag{B-2}$$

is increasing in q at $q = q_g$. Lemma A1 establishes that \bar{n} is truly the marginal user of the public service.

Lemma A1. *The demand for q is monotonically increasing in ability whenever q is a normal good. Specifically, $q_n = q_g$ for $n \leq \bar{n}(q_g)$, and $q_n > q_g$ for $n > \bar{n}(q_g)$.*

Proof. The no bunching condition ensures that every n earns more when choosing private provision than when choosing public provision. Subtracting the right-hand side from the left-hand side of the condition for participation in the public program (which is (B-1) evaluated at n) and differentiating w.r.t. n yields

$$v'(y_{np}/n) \frac{y_{np}}{n^2} - v'(y_{ng}/n) \frac{y_{ng}}{n^2} > 0$$

since $y_{np} > y_{ng}$. Since the participation condition is satisfied with equality for $n = \bar{n}(q_g)$, the result follows. \square

Lemma A1 implies that the government budget constraint may be written as follows

$$R = \int_0^\infty T(y_n) f(n) dn - \int_0^{\bar{n}} pq_g = \int_0^\infty T(y_n) f(n) dn - F(\bar{n}) pq_g \geq 0.$$

The government implements universal public provision of q_g and wishes that all individuals $n < \tilde{n}$ are serviced by the public sector, irrespective of individual preferences. In addition, the government cares about the well-being of all citizens. Let q_g^* be the

optimal level of q_g when the government does not operate subsidies to private purchases of q . There is no need to specify q_g^* more precisely, what matters is that for any choice of q_g^* there is a marginal user, $\bar{n}(q_g^*)$, given by (B-1). Presumably, q_g^* is chosen such that $\bar{n}(q_g^*)$ is a fixed point of the social welfare function that determines the optimal q_g as a function of the marginal user, $\bar{n}(q_g)$. Further, let $q_g^*(\tilde{n})$ be the optimum when q_g is set to serve the interests of $n < \tilde{n}$ only. When $D \equiv \{n \mid \tilde{n} \leq n \leq \bar{n}[q_g^*(\tilde{n})]\} \neq \emptyset$ the problem of public sector dominance exists.

We may now state the main proposition.

Proposition A1. *When the problem of public sector dominance exists, a Pareto improvement arises if the government subsidizes private purchases of q at a rate just sufficient to induce \bar{n} to opt for private provision. Further, under universal provision the optimal policy scheme subsidizes private purchases of q for all individuals with ability $n \leq \bar{n}$ at a decreasing rate and sets $q_g = q_g^*(\tilde{n})$.*

Proof. Let s be a lump sum subsidy on private purchases of x available to all agents and let \bar{n}_s be the marginal user of public provision. By Lemma A1, the effect of s on the government budget is

$$\begin{aligned} \frac{\Delta R}{\Delta s} &= \int_0^\infty \frac{\Delta T(y_n)}{\Delta s} f(n) dn - [F(\bar{n}_s) - F(\bar{n})] pq_g - [1 - F(\bar{n}_s)] s \\ &= \int_{\bar{n}_s}^\infty \frac{\Delta T(y_n)}{\Delta s} f(n) dn - [F(\bar{n}_s) - F(\bar{n})] pq_g - [1 - F(\bar{n}_s)] s. \end{aligned}$$

It follows that a marginal subsidy increases revenue

$$s \longrightarrow 0 \implies \frac{\Delta R}{\Delta s} \longrightarrow \frac{dR}{ds} = [T(y_{np}) - T(y_{ng}) + pq_g] f(\bar{n}) > 0,$$

since $T(y_{np}) \geq T(y_{ng})$ by Lemma A1. This creates scope for a Pareto improvement since the introduction of the subsidy leaves the utilities of all agents unaffected (except for \bar{n} whose utility marginally increases).

The proof now evolves as follows. First, I show that it is optimal to implement type-specific subsidize to private purchases of q for all agents $n \geq \tilde{n}$. Then, I show that the optimal policy can be implemented using only observable earnings.

For any choice of q_g^* , some individuals will satisfy (B-2). Take any such individual n' . If it were possible to provide n' with an idiosyncratic lump sum subsidy on her private purchases of q , the government would wish to do so. The required subsidy $s_{n'}$ satisfies

$$u(z_{n'p} - \pi q_{n'p} + s_{n'}, q_{n'p}) - v(y_{n'p}/n') = u(z_{n'g}, q_g^*) - v(y_{n'g}/n') + \varepsilon \quad (\text{B-3})$$

for some $\varepsilon > 0$. It is easily seen that $q_{n'p} > q_g^*$. Assume that $s_{n'} \geq \pi q_g^*$. In this case, there exists an $\varepsilon > 0$ such that

$$\begin{aligned} u(z_{n'g}, q_g^*) - v(y_{n'g}/n') + \varepsilon &= u(z_{n'p} - \pi q_{n'p} + s_{n'}, q_{n'p}) - v(y_{n'p}/n') \\ &\geq u[z_{n'p} - \pi(q_{n'p} - q_g^*), q_{n'p}] - v(y_{n'p}/n') \\ &\geq u[z_{n'g} - \pi(q_{n'g} - q_g^*), q_{n'g}] - v(y_{n'g}/n') \\ &> u(z_{n'g}, q_g^*) - v(y_{n'g}/n') + \varepsilon \end{aligned}$$

by (B-2). This is a contradiction and, thus, $s_{n'} < \pi q_g^*$ must hold. It follows that a policy of subsidizing the private purchases of q by n' at the rate $s_{n'}$ slightly increases the utility of n' , leaves everyone else unaffected, and saves government funds. Therefore, if feasible, this policy must be part of the optimum. As a result, it is optimal to introduce type-

specific subsidies to all agents that satisfy (B-2) at the initial q_g^* . Since the demand for q is monotonically increasing in ability, the new optimum must have an optimal q_g that is less than q_g^* . Repeating the above argument it is then optimal to introduce type-specific subsidies to all agents that satisfy (B-2) at the new optimum q_g . Iterating the argument shows that it is optimal to operate type-specific subsidies to all agents $n \geq \tilde{n}$.

It is possible to implement the optimal policy by only conditioning subsidies on observable earnings. Indeed, an income-dependent subsidy profile, defined by the sequence of equalities given by (B-3) for $n \geq \tilde{n}$ gives the same marginal utility premium at all income levels and thus does not disturb the truthful revelation. It follows directly from Lemma A1 that s is decreasing in income. Further, $s = 0$ for $n > \bar{n}(\tilde{q}_g)$ where \tilde{q}_g is the optimal value of q_g when only individuals with $n < \tilde{n}$ opt for public provision.

An alternative implementation applies a uniform subsidy to all citizens (or, everyone except $n > \bar{n}$). The optimal subsidy, s^* , is then given by

$$u(z_{\tilde{n}p} - \pi q_{\tilde{n}p} + s_{\tilde{n}}, q_{n'p}) - v(y_{\tilde{n}p}/\tilde{n}) = u(z_{\tilde{n}g}, q_g^*) - v(y_{\tilde{n}g}/\tilde{n}) + \varepsilon.$$

The tax function should then be adjusted to effectively phase-out the excess subsidy. The resulting increase in marginal tax rates simply corrects the excess transfer embodied in the uniform subsidy. It is clear that other implementations are possible such as a stepwise decrease in the subsidy as long as the income tax is adjusted accordingly. The common feature across all implementations is that the effective value of the subsidy decreases with income and equals zero for all income above $y_{\tilde{n}p}$. \square

References

- [1] Blackorby, C. and D. Donaldson (1988). “Cash Versus Kind, Self-Selection, and Efficient Transfers.” *American Economic Review*, 78, pp. 691–700.
- [2] Bertrand, M., S. Mullainathan and E. Shafir (2003). “A Behavioral-Economics View of Poverty.” *American Economic Review (Papers & Proceedings)*, 94, pp. 419–423.
- [3] Besley, T. (1991). “Welfare Improving User Charges for Publicly Provided Private Goods.” *Scandinavian Journal of Economics*, 93, pp. 495–510.
- [4] Besley, T. and S. Coate (1991). “Public Provision of Private Goods and The Redistribution of Income.” *American Economic Review (Paper & Proceedings)*, 81, pp. 979–984.
- [5] Blomquist, S. and V. Christiansen (1995). “Public Provision of Private Goods as a Redistributive Device in an Optimum Income Tax Model.” *Scandinavian Journal of Economics*, 97, pp. 547–567.
- [6] Blomquist, S. and V. Christiansen (1998). “Topping Up or Opting Out? The Optimal Design of Public Provision Schemes.” *International Economic Review*, 39, pp. 399–411.
- [7] Boadway, R. and M. Marchand (1995). “The Use of Public Expenditures for Redistributive Purposes.” *Oxford Economic Papers*, 47, pp. 45–59.

- [8] Diamond, P.A. (2003). *Taxation, Incomplete Markets, and Social Security: the 2000 Munich Lectures*. Cambridge: MIT Press.
- [9] Diamond, P.A. (2006). “Optimal tax treatment of private contributions for public goods with and without warm glow preferences.” *Journal of Public Economics*, 90, 897–919.
- [10] Feldstein, M. (2005). “Rethinking Social Insurance.” *American Economic Review*, 95, pp. 1–24.
- [11] Golosov, M. and A. Tsyvinski (2006). “Designing Optimal Disability Insurance: A Case for Asset Testing.” *Journal of Political Economy*, 114, pp. 257–269.
- [12] Musgrave, R.M. (1959). *The Theory of Public Finance*. New York: McGraw-Hill.
- [13] Nichols, A.L. and R.J. Zeckhauser (1982). “Targeting Transfers through Restrictions on Recipients.” *American Economic Review (Papers & Proceedings)*, 72, pp. 301–317.
- [14] Rothschild, M. and J. Stiglitz (1976). “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” *Quarterly Journal of Economics*, 90, pp. 629–649.
- [15] Stiglitz, J. (1982). “Self-Selection and Pareto Efficient Taxation.” *Journal of Public Economics*, 17, pp. 213–240.

Chapter 3

Optimal Provision of Public Goods: A Synthesis

Claus Thustrup Kreiner
University of Copenhagen, EPRU, and CESifo

Nicolaj Verdelin
University of Copenhagen and EPRU

August 2008

Abstract

This paper presents a general framework to analyze the optimal provision of public goods. There currently exist two competing approaches in the literature. The traditional approach highlights the importance of distortionary taxation and distributional concerns. The new approach neutralizes distributional concerns by adjusting the non-linear income tax, and finds that this reinvigorates the simple Samuelson rule when preferences are separable in goods and leisure. We provide a synthesis by demonstrating that both approaches derive from the same basic formula. Using the principles of the new approach, we derive a fully general, intuitive formula for the optimal level of a public good. This formula shows that distortionary taxation may have a role to play as in the traditional approach. However, the main determinants of optimal provision are completely different and the traditional formula with its emphasis on MCF only obtains in a very special case.

Keywords: Public goods; Samuelson rule, Social marginal cost of public funds, Benefit principle

JEL classification: H41; H23; H11

1 Introduction

The practice of cost-benefit analysis is hugely influential in everyday government decisionmaking on public projects throughout the developed world. According to the Danish Ministry of Finance, the evaluation of public projects in Denmark assumes that the cost of financing is 1.2 times the actual expenditure (corresponding to the official Danish Marginal cost of funds, MCF). Similarly, a CBPP report (Stone et al., 2008) explains that the Congressional Budget Office assumes that 25% of the proceeds from environmental taxes are lost due to the impact of behavioral responses on revenues from, e.g., the income tax. This corresponds to an MCF of about 1.3. While it is common practice to assume that the financing of public projects involves an excess burden as captured by the MCF, the theoretical foundation for such a practice is much less clear.

Historically, the study of the optimal supply of public goods has formed three ‘waves’ (Kaplow, 2004) differing by the underlying assumptions. The first wave originates from Samuelson (1954), which established the famous Samuelson rule that equates the sum of the marginal willingness to pay for the public good of all citizens to the marginal rate of transformation (MRT). This result is derived in a first best setting where individual lump sum taxes are available. Atkinson and Stern (1974) later criticized the Samuelson rule on the grounds that the first best is not attainable. Instead, their analysis relies on distortionary taxation and arrives at an extended Samuelson rule where the effective cost of public goods is identified as MCF times MRT. This approach has been very influential and also underlies the survey of Ballard and Fullerton (1992).

The second wave approach has since been further developed to allow for heterogeneity in earnings abilities across households (Dahlby, 1998). This has led Slemrod and Yitzhaki (2001) and Gahvari (2006) to argue that the evaluation of public projects should take account, not only of MCF, but also of the marginal benefit of projects. Thus, distributional concerns are shown to matter for the optimal level of public goods.

In contrast, the third wave of the theory of public goods argues that distributional concerns are irrelevant to the evaluation of public projects. This line of research, initiated by Hylland and Zeckhauser (1979) and further pursued by Christiansen (1981) and Kaplow (1996), holds that unintended distributional effects can be undone by the income tax. Formally, Christiansen (1981), in the context of the optimal non-linear income tax, and Kaplow (1996), more generally, show that the original Samuelson rule applies when preferences are separable in leisure and other goods, including the public good. However, less effort has been devoted to the study of the optimal policy rule when the separability assumption fails.

The divergent results have created a state of confusion as illustrated by the debate in the wake of Kaplow's (2004) survey (see Goulder et al., 2005, and the reply by Kaplow). The recent developments in the second wave have directed attention to cases when the weak separability assumption is not fulfilled but have, at the same time, made a departure from the *benefit principle* that underlies the third wave. The benefit principle builds on the flexibility of the non-linear income tax and holds that each individual should contribute to the financing of the public good corresponding to her own marginal willingness

to pay. This eliminates any distributional concerns due to the specifics of the financing scheme and instead focuses on the public goods problem itself. By deviating from this principle, it becomes unclear whether the results of the second wave are driven by genuine distributional concerns or whether they are due to the assumed financing scheme.

This paper uses a completely general framework to reconcile the results of the second and third waves. Specifically, we demonstrate that both approaches derive from the same general formula requiring that a public project is completed only when the social marginal benefit of the project (SMB) exceeds the social marginal cost of public funds (SMCF). However, a discrepancy arises between the two waves because the second wave does not impose any systematic restrictions on the financing scheme underlying the public project. Our analysis begins by generalizing the fundamental result of the second wave (Gahvari, 2006) and then applies a dual approach to show that the third wave originates from the same formula.

We further apply the benefit principle to analyze the optimal provision of public goods in a fully general model of earnings and preference heterogeneity. We derive an intuitive formula for the optimal level of a public good that identifies the correlation between ability and the marginal willingness to pay for the public good as the driving force behind any deviations from the Samuelson rule. By applying the benefit principle the policy rule derived in this paper eliminates any distributional concerns related to the financing scheme rather than to the public goods problem itself.

Using our approach, we can generalize the result of Kaplow (1996) to functional forms

that allow income to directly affect the willingness to pay for the public good. Indeed, a correlation between the marginal willingness to pay and income has vastly different implications for the optimal level of a public good than does correlations with ability, although the two are observationally equivalent. Only the latter leads to a departure from the Samuelson rule since correlations with income can be made distributionally neutral through appropriate adjustments of the income tax.

We finally show that there is equivalence between the results of the second and third waves only in a special case when the willingness to pay for the public good is linear in ability. In this case, the basic second wave formula with its emphasis on the MCF obtains even when marginal tax changes are determined by the benefit principle.

The paper is organized as follows. Section 2 presents our basic model with a continuum of agents and preference heterogeneity. Section 3 derives a general formula for the optimal level of a public good when there are no restrictions on the financing scheme as in the second wave. Section 4 shows the relationship between the second and third wave and derives a general, intuitive formula for the optimal level of a public good when marginal tax changes are governed by the benefit principle. In Section 5 we provide a special case when the two approaches lead to identical results, and the standard second wave formula with its emphasis on MCF applies. Finally, Section 6 offers a few concluding remarks.

2 The Model

This section presents a fully general framework to analyze the optimal provision of public goods. The model has a continuum of agents, each characterized by an innate ability n , which is also our index of identification. The distribution of abilities across the population is given by the non-degenerate density function $f(n)$. Consumers allocate their income to consumption c , which could be thought of as either a vector of consumption goods or a single composite good. Gross earnings or, more generally, taxable income is denoted z . The utility of agent n is given by

$$u(c, g, z, n), \tag{1}$$

where $u_c = \partial u / \partial c > 0$, $u_g > 0$, $u_z < 0$, and $u(\cdot)$ is quasiconcave. This utility specification embodies preference heterogeneity across individuals of different abilities. It also encompasses the traditional Mirrleesian specification, $u(c, g, z, n)$, as a special case. The term z/n builds on the notion that more able persons must exert less effort to attain a given income level. If this logic is extended to other domains of everyday life, as in Becker (1965), it seems natural that ability should have an impact on the utility of consuming, as long as the skills of home production are correlated with market productivity. The theory of household production views market goods as an input in a production process, which, along with individual skills, determines the output that ultimately enters individual utility. Thus, persons of different skills may benefit differently from a given input of c or g . For instance, an individual's ability to cook determines the utility derived

from a basket of groceries. Similarly, the utility derived from public goods such as the police or the judicial system depends on both the skill and need to benefit from such institutions, which is likely influenced by individual ability. Thus, the formulation in (1) captures both innate preference differences between individuals of different abilities and preference differences due to the technology of home production.

Since the government cannot condition taxes on the unobservable ability it is forced to operate a (possibly) non-linear income tax function $T(z, \theta)$, where θ is a shift parameter used to capture the effects of changes to the tax function. The budget constraint of agent n is

$$c \leq z - T(z, \theta) + I(n), \quad (2)$$

where $I(\cdot)$ is non-labor (net) income. The first-order conditions for the choice of c and z imply

$$\text{MRS}_{cz} \equiv -\frac{u'_z(\cdot)}{u'_c(\cdot)} = 1 - m, \quad (3)$$

where $m \equiv \partial T(\cdot) / \partial z$ is the marginal tax rate at the income level z . The earnings choice of the household may be written as $z = z((1 - m), y, g)$, where $(1 - m)$ is the marginal net-of-tax rate and $y \equiv mz - T(z, \theta) + I$ is virtual income.

The indirect utility function is $v(n) \equiv u(c(n), g, z(n) / n)$ and gives the utility level of individual n when consumption and labor supply are chosen optimally. We follow the standard approach in optimal taxation and contract theory and assume that (i) utility is increasing in ability, i.e., $\partial u / \partial n > 0$, and (ii) the Spence-Mirrlees single-crossing

condition is satisfied (e.g., Salanié, 2003):

$$\text{MRS}_{cz} \text{ is decreasing in } n. \tag{4}$$

The first assumption along with the Envelope Theorem ensures that the indirect utility is increasing in ability, $dv/dn = \partial u/\partial n > 0$. The second assumption ensures that the tax system is implementable, i.e., that higher ability individuals always choose higher equilibrium earnings, implying that the government can use income as a signal of the underlying ability.

The government cares about redistribution as well as the provision of public goods. The preferences of the government are captured by a Bergson-Samuelson social welfare function of the form

$$\Omega = \int_n \Psi(u(\cdot)) f(n) dn, \tag{5}$$

where $\Psi(\cdot)$ is a concave function reflecting the distributional concerns of the policymaker. The cost of producing g units of the public good in units of the numeraire is normalized to g , without any loss of generality. The government budget constraint is

$$R \equiv \int_n T(z, \theta) f(n) dn - g(\theta) \geq 0,$$

where the public goods nature of g is seen from the fact that g enters only once in the government budget constraint but still appears in everyone's utility functions.

Changes to the reform parameter θ capture both the changes in the supply of g and the associated adjustment of the tax function. Differentiating (5) using the first-order

condition (3) yields the effect of a marginal reform, $d\theta$, on social welfare

$$\frac{d\Omega}{d\theta} \frac{1}{\lambda} = - \int_n \omega(n) \frac{\partial T(z, \theta)}{\partial \theta} f(n) dn + \frac{dg}{d\theta} \cdot \int_n \omega(n) \frac{u_g}{u_c} f(n) dn, \quad (6)$$

where $\lambda \equiv \int_n \Psi'(\cdot) u_c(\cdot) f(n) dn$ and $\omega(n) \equiv \frac{\Psi'(u(\cdot)) u_c(\cdot)}{\lambda}$ denotes the social welfare weight of agent n . Similarly, the effect of a reform on government revenue is given by

$$\frac{dR}{d\theta} = \int_n \left(\frac{\partial T(z, \theta)}{\partial \theta} + m \frac{dz}{d\theta} \right) f(n) dn - \frac{dg}{d\theta}, \quad (7)$$

where the first term under the integral is the direct revenue effects and the second term captures the effect of behavioral responses on government revenue. These behavioral responses are driven by changes to the tax schedule as well as by effects of government consumption on household utility.

3 Second Wave

The second wave view of optimal public goods supply is due originally to Atkinson and Stern (1974) and has exerted a tremendous influence on the practice of cost-benefit analysis (e.g., Ballard and Fullerton, 1992). This approach to deriving a formula for the optimal public goods supply does not impose any restrictions on the financing scheme other than the requirement that the reform is fully financed, i.e., $dR/d\theta = 0$. From (7) this yields

$$\frac{dg}{d\theta} = \frac{\int_n \left(\frac{\partial T(z, \theta)}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial(1-m)} \frac{\partial m}{\partial \theta} \right) \right) f(n) dn}{1 - \int_n m \frac{\partial z}{\partial g} f(n) dn},$$

where the behavioral responses to taxable income have been decomposed into tax effects and an effect from the change in g .¹ A marginal expansion of g is desirable if it increases

¹When deriving the behavioral responses to the tax reform, we follow the second wave literature and assume that the tax schedule is piece-wise linear. This ensures that there is no feed-back effect from the

social welfare, i.e., if $d\Omega/d\theta \geq 0$. Insert the above expression in (6) and apply this test to get

$$\frac{\int_n \omega(n) \frac{u_g}{u_c} f(n) dn}{1 - \int_n m \frac{\partial z}{\partial g} f(n) dn} \geq \frac{\int_n \omega(n) \frac{\partial T}{\partial \theta} f(n) dn}{\int_n \left[\frac{\partial T}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial(1-m)} \frac{\partial m}{\partial \theta} \right) \right] f(n) dn}. \quad (8)$$

The uncompensated elasticity of taxable income with respect to the net-of-tax rate is defined as $\varepsilon \equiv \frac{1-m}{z} \frac{\partial z}{\partial(1-m)}$. From the Slutsky-equation, it may be decomposed into a compensated elasticity and an income effect, that is $\varepsilon = \varepsilon^c - \eta$ where ε^c is the compensated elasticity and $\eta \equiv -(1-m) \frac{\partial z}{\partial I}$ is the income effect.² Further, let

$$\Phi \equiv \frac{\partial m / \partial \theta}{\partial a / \partial \theta}, \quad s(n) \equiv \frac{\partial T}{\partial \theta} / \left(\int_n \frac{\partial T}{\partial \theta} f(n) dn \right),$$

where a is the average tax rate. The parameter Φ captures the progressivity of the implied tax reform, and $s(n)$ is the share of the direct tax changes that is borne by agent n . Using this we can rewrite (8) in terms of behavioral elasticities to arrive at Proposition 1.

Proposition 1 *A marginal expansion of a public good is desirable iff*

$$\frac{\int_n \omega(n) MRS_{cg} f(n) dn}{1 - \int_n m \frac{\partial z}{\partial g} f(n) dn} \geq \frac{\int_n \omega(n) s(n) f(n) dn}{\int_n \left(1 - \frac{m}{1-m} (\Phi \cdot \varepsilon^c - \eta) \right) s(w) f(n) dn}. \quad (9)$$

change in z to the marginal tax rate, and thus no additional earnings responses beyond those triggered directly by the tax reform. Mathematically, we avoid including second derivatives of the tax function into the formula T'' . The assumption of piece-wise linearity implies that there will be bunching at the various kinks in the tax schedule. This does not constitute a problem for our final results but may imply that taxable income elasticities are zero at a kink point because marginal changes are not sufficient to move the individual away from the kink point.

²Previous contributions have defined hours-of-work elasticities. The elasticity of taxable income captures hours-of-work responses as well as all other behavioral responses that are relevant for total tax payments, and the empirical evidence indicates that this elasticity is significantly larger than the hours-of-work elasticity (e.g. Gruber and Saez, 2002).

Proof. See Appendix A. ■

Expression (9) gives a completely general formula for the optimal g and generalizes the result in Gahvari (2006) to a continuous setting with a more general utility specification. Intuitively, a marginal expansion of the public good is desirable when the social marginal benefit (SMB, the left-hand side) of the public good exceeds the distribution-weighted social marginal cost of public funds (SMCF, the right-hand side), i.e., $SMB \geq SMCF$. Indeed, the right-hand side is the continuous-setting equivalent to the social marginal cost of public funds derived in Dahlby (1998).³

Proposition 1 shows that distributional considerations affect the optimal level of a public good. This includes both the distribution of the benefits from an expanded g and the distribution of the costs of financing the expansion. However, condition (9) suffers from the absence of any restrictions on the financing scheme. Indeed, the second wave does not impose any systematic relationship between individual benefits from g and individual contributions to the financing of the expanded g . In principle, the marginal tax changes implied by the reform could follow any arbitrary pattern. This ignores the flexibility of the non-linear income tax, and thereby assigns a role to distributional considerations that are unrelated to the problem of public goods provision (see also Auerbach and Hines, 2002). Instead, distributional consequences of the specific financing scheme affect the optimal g . This approach may have merit when there are exogenous

³A similar formula for the marginal cost of public funds appears in Kleven and Kreiner (2006), who include extensive responses as well. We have chosen to follow the tradition in analyses of the optimal provision of public goods and MCF by focusing on intensive responses only.

constraints that limit the adjustment of the tax schedule as emphasized by Slemrod and Yitzhaki (2001) and Gahvari (2006), but is not a useful principle for general normative prescriptions. The latter should reflect all available instruments as is the case with analyses of the optimal non-linear income tax. As an example, consider a public good that yields uniform benefits to all individuals. Now suppose that the financing scheme is proportional and simply adjusts everyone's marginal tax rate slightly upward to finance the public good. In this case, the policy rule (9) might reject the public project even when the sum of the benefits exceeds the direct cost because the financing scheme introduces unnecessary distortions of the earnings decision. If we instead used our knowledge of the distribution of the benefits of g , we would choose lump sum financing, avoid any distortions, and approve the project whenever the Samuelson rule is satisfied.

4 Third Wave

This section takes the analysis of public goods provision a step further by imposing a restriction on the relationship between the benefits of an expanded g and the individual marginal tax changes implied by the associated tax reform. Thus, the change to the entire tax schedule is endogenous, i.e., at every income level both the direct change to the tax burden and the change in the marginal tax rate are determined endogenously by the individual benefits from the reform. This approach, which we labelled the third wave in the Introduction, keeps individual utilities unchanged. As a result, we cannot apply the same method to derive the optimal level of g as in the second wave approach

of the previous section. Instead, we use an alternative approach that keeps social welfare unaffected and determines the desirability of a marginal expansion of g by calculating the effect of the reform on government revenue. If the effect is positive, the reform is socially desirable. We show in Appendix B that setting $d\Omega/d\theta = 0$ implies that $dR/d\theta \geq 0$ is equivalent to

$$\frac{\int_n \omega(w) \frac{u_g}{u_c} f(n) dn}{1 - \int_n m \frac{\partial z}{\partial g} f(n) dn} \geq \frac{\int_n \omega(w) \frac{\partial T}{\partial \theta} f(n) dn}{\int_n \left[\frac{\partial T}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial(1-m)} \frac{\partial m}{\partial \theta} \right) \right] f(n) dn}, \quad (10)$$

which is identical to (8). The fact that we arrive at the same formula as in the second wave is not surprising since we have merely applied the dual approach to determine the optimal level of g . Importantly, the equivalence of (8) and (10) provides a link between the second and third wave. Indeed, both approaches derive from the same basic formula. The differences lie in the assumptions made regarding the associated tax reform.

Unlike the second wave, which does not impose any restrictions on the financing scheme other than budget balance, the procedure advanced in this paper evaluates the benefits of an expansion of the public good by use of the *benefit principle*, introduced by Hylland and Zeckhauser (1979) and applied by Christiansen (1981) and Kaplow (1996, 2004). According to this principle, a (marginal) expansion of g should be financed by a benefit-offsetting, or distribution neutral, change in the tax function. Since the reform keeps individual utilities unaffected the merits of a marginal expansion of g depend on the implied changes to government revenue, i.e., if $dR/d\theta \geq 0$ the expansion of g should be implemented. Thus, we consider a reform, $d\theta$, that affects g as well as the tax function

$T(\cdot)$ such that

$$\frac{dv(n)}{d\theta} = 0 \text{ for all } n, \quad (11)$$

$$\frac{dv'(n)}{d\theta} = 0 \text{ for all } n. \quad (12)$$

The benefit-offsetting expansion of g adjusts the tax function to capture the benefits of the additional g from each individual n . Since the tax function depends on income, not ability, the reform may have distortionary effects on the incentive to work.

It follows from (11) and the envelope theorem that

$$\frac{\partial T(z, \theta)}{\partial \theta} = \frac{u'_g(\cdot)}{u'_c(\cdot)} \cdot \frac{dg}{d\theta} = \text{MRS}_{cg} \cdot \frac{dg}{d\theta}. \quad (13)$$

Further, differentiating (12) and using the budget constraint gives

$$\frac{dz}{d\theta} = \frac{u''_{cn}(\cdot) \frac{\partial T(\cdot)}{\partial \theta} - u''_{gn}(\cdot) \frac{dg}{d\theta}}{u''_{cn}(\cdot)(1-m) + u''_{zn}(\cdot)}. \quad (14)$$

For any given individual n , the first term in the numerator captures the cost of the increased tax burden, while the second term captures the benefits of the expanded g . Any discrepancy between individual costs and individual benefits functions just like a change in the marginal tax rate and thus affects earnings. Only when costs and benefits are exactly aligned is there no change in individual income since in this case the incentive to supply earnings remains unchanged. This is entirely consistent with the benefit principle, which cannot condition reform changes on the unobservable ability. Generally speaking, (14) shows that we cannot separately evaluate the costs and benefits of a public program. Indeed, it is the individual *net benefit* that determines any income responses.

We show in Appendix B that (14) may be rewritten to

$$\frac{dz}{d\theta} = \frac{\partial MRS_{cg}(c, g, z) / \partial n}{\partial MRS_{cz}(c, g, z) / \partial n} \cdot \frac{dg}{d\theta}, \quad (15)$$

where the single-crossing condition implies that the denominator is negative. The partial derivatives in this expression measure the effect of ability on the marginal rates of substitution between, respectively, c and g in the numerator and c and z in the denominator for given c , g and z .

The application of the benefit principle implies that the expansion of g and the accompanying change in the tax function keeps everyone's utility, and thus social welfare, unchanged. We may rewrite (10) as

$$\frac{\int_n \omega(w) \frac{u_g}{u_c} \frac{dg}{d\theta} f(n) dn}{\frac{dg}{d\theta} - \int_n m \frac{\partial z}{\partial g} \frac{dg}{d\theta} f(n) dn} \geq \frac{\int_n \omega(w) \frac{\partial T}{\partial \theta} f(n) dn}{\int_n \left[\frac{\partial T}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial(1-m)} \frac{\partial m}{\partial \theta} \right) \right] f(n) dn}.$$

Now (13) gives $\int_n \omega(w) \frac{u_g}{u_c} \frac{dg}{d\theta} f(n) dn = \int_n \omega(w) \frac{\partial T}{\partial \theta} f(n) dn$ implying

$$\int_n \left[\frac{\partial T}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial(1-m)} \frac{\partial m}{\partial \theta} + \frac{\partial z}{\partial g} \frac{dg}{d\theta} \right) \right] f(n) dn \geq \frac{dg}{d\theta},$$

where the second term in the parenthesis under the integral is simply $dz/d\theta$. By inserting (13) and (15), we get Proposition 2.

Proposition 2 *A marginal expansion of a public good is desirable iff*

$$\int \left(MRS_{cg} + m \cdot \frac{\partial MRS_{cg} / \partial n}{\partial MRS_{cz} / \partial n} \right) f(n) dn \geq MRT. \quad (16)$$

Proposition 2 shows that the Samuelson rule must be amended by a term that is affected by the *partial* correlation between ability and the marginal willingness to pay for the

public good. The additional term corrects for the revenue implications of the behavioral responses to the reform. The optimal level of g is affected by correlations with the unobservable n because the tax function is constrained to depend on the imperfect signal that is income. Generally, the use of any government instrument that indirectly redistributes based on ability should reflect the desire to redistribute income in the most efficient way. Thus, the optimal g is adjusted according to who benefits the most from the public good as long as the differences in benefits are due to differences in ability. This result does not require that the income tax is optimized, only that the associated tax reform is guided by the benefit principle. It is important to note that the partial effects on the MRS's in (16) are evaluated at a given income level. Thus, variations in MRS due entirely to variations in z do not affect the optimal public goods supply. Rather, the crucial test is whether the slope of the indifference curves of people of different ability differ *when evaluated at the same consumption-income bundle*. If this is the case, the public good effectively redistributes based on the unobservable ability.

Intuitively, when marginal tax rates are positive, the supply of public goods is reduced relative to the first best if the marginal willingness to pay for the public good increases with ability. In this case, the benefit principle implies that higher incomes must contribute more to the financing of the public good. However, part (or all) of the additional benefit enjoyed by persons with higher incomes stems from their innate ability and is realized independently of the chosen income level. Thus, the additional taxes implied by the reform reduce the incentive to work. The size of the additional distortion depends on

the responsiveness of earned income as captured by the denominator of the second term above. Also, the stronger is the influence of ability on the marginal willingness to pay, the more difficult it is for the government to finance g in a non-distortionary fashion. An alternative way to view this result focuses on how the concern for redistribution affects the optimal level of g . When persons of higher ability benefit relatively more from the presence of the public good, the supply of g adversely affects the government's scope for redistribution. Indeed, the public good effectively redistributes in favor of the rich. This point applies the same logic as do Nichols and Zeckhauser (1982) and Blackorby and Donaldson (1988) in the context of in kind transfers. Also, Kaplow (2008) provides a similar intuition for the case of public goods but does not arrive at our general formula (16). A reversal of this argument explains why the public goods supply is *higher than in the first best* when there is a negative correlation between ability and the marginal willingness to pay for the public good. In this case, supplying g provides an additional means to redistribute in favor of the poor.

A theater performance is an example of a public good that may be valued higher by the more able for a given income. If so, the optimal public financial support for theaters is less than the Samuelson rule predicts because such support effectively redistributes income towards the more able. In contrast, public transportation is likely to benefit persons of lower ability more for a given income. Efficient public transportation reduces the travel time to and from the workplace, leaving more time for other activities. A low ability individual must work longer hours to uphold a given income and therefore, pre-

sumably, values her sparetime more. Thus, subsidies to public transportation effectively redistribute income towards the less able, over and above what is attainable through the income tax. Importantly, consumption patterns across incomes do not necessarily reveal the desirability of public transport subsidies. If low income individuals choose public transport because they cannot afford a car, not because they are of low skill, the Samuelson rule still applies.

Proposition 2 also clarifies when the first best obtains. The sufficient condition is that there is no partial effect from ability to the willingness to pay for the public good. Remember that the partial effects from ability to MRS in (16) are evaluated at a given income level. Thus, the crucial question for the determination of the optimal g is whether the marginal willingness to pay is different for a person of high ability when she imitates the choices of a lower ability individual. If this is not the case, and people of different ability have the same MRS_{cg} for given z , the Samuelson rule applies and distributional considerations should not affect the level of the public good. This does not rule out that people of different ability, as they position themselves at different income levels, have different willingnesses to pay *in equilibrium*. In this case, the financing of the public good is not uniform under the benefit principle and, as a result, marginal tax rates are affected. But these tax variations are not distortionary as the marginal willingness to pay also varies with income. Differential financing is only distortionary when taxpayers can avoid the additional burden without reducing the benefit they enjoy from the public good. Thus, armed with Proposition 2 we can generalize the result of Kaplow (1996) to

a more general utility function.⁴

Corollary 1 *Assume that individual utility is separable:*

$$u(c, g, z, n) = \tilde{u} [w^1(c, g, z), w^2(z, n)].$$

Then an expansion of g is socially desirable whenever the Samuelson condition holds, i.e.,

$$\int MRS_{cgf}(n) dn \geq MRT.$$

Proof. See Appendix B. ■

The above utility specification implies that variations in the marginal willingness to pay for the public good derive from income directly, not the underlying ability. If the marginal willingness to pay increases with income, the benefit principle implies that marginal tax rates increase as a result of the reform but these changes are not distortionary as the individual benefit from the public good also increases with income (see Blomquist et al., 2008, for a similar point).

When utility is given by the standard Mirrleesian specification $u(c, g, z/n)$ the formula for the optimal g can be written in terms of correlations between the marginal willingness to pay for g and labor supply, $l = z/n$. This is because with the standard utility function any correlation with n implies a correlation with l .

⁴Observe that a comparison of the second wave formula (9) to the Samuelson rule reveals that when preferences are weakly separable in ability and the consumption of c and g the tax reform that results from the benefit principle implies that

$$\frac{u_g}{u_c} \frac{1}{1-m} (\Phi \cdot \varepsilon^c - \eta) + \frac{\partial z}{\partial g} = 0.$$

Corollary 2 *When individual utility is given by*

$$u(c, g, z, n) = \tilde{u}(c, g, z/n)$$

an expansion of g is socially desirable whenever

$$\int \left(MRS_{cg} + m \frac{\partial MRS_{cg}/\partial l}{\partial MRS_{cz}/\partial l} \right) f(n) dn \geq MRT.$$

Proof. See Appendix B. ■

When ability is restricted to affect utility only through l , the evaluation of a public project departs from the Samuelson rule if the marginal willingness to pay for the public good depends on individual labor supply. Thus, if MRS_{cg} decreases with l the optimal level of the public good is less than predicted by the Samuelson rule. In this case, the public good is valued relatively more by those who must deliver fewer working hours to attain a given income, i.e., people of higher ability. Therefore, the public good impacts negatively on the government's ability to redistribute income. However, the opposite situation is equally plausible. When MRS_{cg} increases with l the optimal g is higher than the first best level. But remember that the importance of correlations with l is merely an artifact of the special shape of the utility function.

5 An Example of Equivalence Between The Two Waves

Generally, the formula for the optimal g deviates from Proposition 2 when the associated tax reform is not governed by the benefit principle. Thus, the second wave approach generally leads to different results than those obtained in the previous section. However,

in one special case the two approaches are equivalent and the simplest form of the second wave formula obtains. The latter holds that public goods should be expanded if

$$\int \text{MRS}_{cg} f(n) dn \geq \text{MCF} \cdot \text{MRT},$$

where MCF is the marginal cost of public funds. In fact, the role assigned to MCF in the determination of the optimal g is somewhat misleading, since the determination of MCF builds on a conceptual experiment that ignores the purpose of the tax change. In contrast, according to the benefit principle, the implied tax change is closely linked to the use of public funds. The discrepancy between these two approaches stands out most clearly when ability does not affect the marginal willingness to pay for the public good. As we have seen, the implied tax changes have no distortionary effects in this case, and therefore do not affect the optimal level of g . In contrast, the same tax reform would obviously distort earnings if the additional revenue was not put to valuable use. As it turns out, the simple correction for the distortionary effects of taxation obtains only in a special case.

Assume utility is given by

$$u = c + nw(g, z) - n \cdot h(z/n),$$

where the functional form of the disutility of labor is taken from Saez (2001) and implies that n reflects potential earnings. We show in Appendix C that a marginal expansion of g is desirable iff

$$\int u'_g(\cdot) \left(1 - \frac{m}{1-m}\varepsilon\right) f(n) dn \geq \text{MRT}, \quad (17)$$

which identifies MCF as a central determinant of the optimal g . If, in addition, the income tax is linear and the elasticity of taxable income is constant across individuals, the condition simplifies to

$$\int \text{MRS}_f(n) \, dn \geq \frac{1}{1 - \frac{m}{1-m}\varepsilon} \cdot \text{MRT} = \text{MCF} \cdot \text{MRT},$$

which is identical to the simple second wave formula (Browning, 1987, Dahlby, 1998, and Ballard and Fullerton, 1992). Thus, only when utility from the public good is linear in ability and the tax system is proportional is the traditional MCF correction valid.

6 Concluding Remarks

The central challenge involved in decisions on the optimal level of a public good is that correlations between the marginal willingness to pay and, respectively, ability and income are observationally equivalent but have vastly different policy implications as first noted by Hylland and Zeckhauser (1979). For instance, are wealthy people overrepresented among opera audiences because they are wealthy, or because they are, generally, better educated? For some purposes casual observation may be sufficient to decide on the desirability of a public project. When more detailed analyses are called for, one is left to search for observable characteristics that have a known (or estimable) relationship with ability.

While the analysis in this paper has focused on public goods, the results may be directly applied to the correction of externalities. We may think of g as a global externality and MRS_{cg} as the willingness to pay for a marginal reduction of the externality. The cost

of reducing g is then the costs of, e.g., abatement or alternative production methods. As argued by Kaplow and Shavell (2002), the most efficient way to regulate externalities is through a price scheme that reflects marginal harm. When consumption patterns differ across individuals, the costs and benefits of such a scheme may be unevenly distributed. However, any distributional effects that are driven by preference variations due directly to income can be undone through adjustments of the income tax (see also Kaplow, 2006). Only when the willingness to pay for harm reduction is correlated with ability should the externality correction depart from first best rules.⁵

⁵If the externality is not global but affects only part of the population, it is necessary for the results that the income tax can follow the same demographic patterns. For instance, pollution in a major city mainly affects its citizens and compensation schemes must then be designed to affect only the citizens of that same city. This is possible if regional taxes are in place and can be adjusted freely. However, local tax functions are often subject to constitutional restrictions. In this case, and when the externality affects subsets of the population that cannot be explicitly targeted, the benefit principle can no longer be applied and alternative methods must be used.

Appendix A

Proof of Proposition 1: The effect of the tax changes on government revenue is

$$\frac{\partial R}{\partial \theta} \Big|_{g=\bar{g}} = \int_n \left[\frac{\partial T}{\partial \theta} + m \left(\frac{\partial z}{\partial y} \frac{\partial y}{\partial \theta} - \frac{\partial z}{\partial (1-m)} \frac{\partial m}{\partial \theta} \right) \right] f(n) dn,$$

which is identical to the numerator on the right-hand side of (8). The change in virtual income is

$$\frac{\partial y}{\partial \theta} = z \frac{\partial m}{\partial \theta} + m \frac{\partial z}{\partial \theta} - m \frac{\partial z}{\partial \theta} - \frac{\partial T(z, \theta)}{\partial \theta} = z \left(\frac{\partial m}{\partial \theta} - \frac{\partial a}{\partial \theta} \right),$$

where $a \equiv T(z, \theta) / z$ is the average tax rate and $\partial a / \partial \theta \equiv \frac{\partial T(z, \theta)}{\partial \theta} / z$. This implies that

$$\frac{\partial z}{\partial \theta} \Big|_{g=\bar{g}} = \left[\frac{\partial z}{\partial y} \left(\frac{\partial m}{\partial \theta} - \frac{\partial a}{\partial \theta} \right) - \frac{1}{1-m} \varepsilon \frac{\partial m}{\partial \theta} \right] z,$$

where $\varepsilon \equiv \frac{1-m}{z} \frac{\partial z}{\partial (1-m)}$ is the uncompensated elasticity of taxable income. We may rewrite

this using the Slutsky equation

$$\frac{\partial z}{\partial \theta} \Big|_{g=\bar{g}} = \left(\eta \frac{\partial a}{\partial \theta} - \varepsilon^c \frac{\partial m}{\partial \theta} \right) \frac{1}{1-m} z,$$

which implies

$$\frac{\partial R}{\partial \theta} \Big|_{g=\bar{g}} = \int_n \left[1 + \frac{m}{1-m} (\eta - \Phi \varepsilon^c) \right] \frac{\partial T}{\partial \theta} f(n) dn,$$

where $\Phi = \frac{\partial m / \partial \theta}{\partial a / \partial \theta}$. Insert this in (8) and use

$$s(n) \equiv \frac{\partial T}{\partial \theta} / \left(\int_n \frac{\partial T}{\partial \theta} f(n) dn \right)$$

to obtain formula (9) in Proposition 1.

A Appendix B

This appendix describes the steps needed to arrive at the equation (15) in Section 4, which is central to Proposition 2. The earnings response is derived from (12), which we repeat here for convenience

$$\frac{dv'(n)}{d\theta} = 0,$$

where $v'(n) = u'_n(c, g, z, n)$. This gives

$$-u_{nc}(\cdot) \frac{dc}{d\theta} - u_{ng}(\cdot) \frac{dg}{d\theta} - u_{nz}(\cdot) \frac{dz}{d\theta} = 0$$

Using $dc/d\theta = (1 - m) dz/d\theta - \partial T(\cdot)/\partial\theta$ we can solve for $dz/d\theta$ above

$$\frac{dz}{d\theta} = \frac{u''_{cn}(\cdot) \frac{\partial T(\cdot)}{\partial\theta} - u''_{gn}(\cdot) \frac{dg}{d\theta}}{u''_{cn}(\cdot) (1 - m) + u''_{zn}(\cdot)},$$

which is identical to (14). Consider the MRS between c and g

$$\text{MRS}_{cg} \equiv \frac{u'_g}{u'_c} \implies \frac{\partial \text{MRS}_{cg}}{\partial n} = \frac{1}{u'_c} \left(u''_{gn} - \frac{u'_g}{u'_c} u''_{cn} \right),$$

and that between c and z

$$\text{MRS}_{cz} \equiv -\frac{u'_z}{u'_c} \implies \frac{\partial \text{MRS}_{cz}}{\partial n} = -\frac{1}{u'_c} (u''_{zn} + (1 - m) u''_{cn}),$$

where we have used (3). Insert this above and use (13) to get

$$\frac{dz}{d\theta} = \frac{\partial \text{MRS}_{cg}/\partial n}{\partial \text{MRS}_{cz}/\partial n} \cdot \frac{dg}{d\theta}.$$

When utility is of the traditional Mirrlessian form, $u(c, g, z/n)$, we can use the relation $z = n \cdot l$ to express the change in z as a function of the dependence of MRS on l instead.

Indeed,

$$\frac{\partial \text{MRS}}{\partial l} = \frac{\partial \text{MRS}}{\partial n} \frac{\partial n}{\partial l} \implies \frac{dz}{d\theta} = \frac{\partial \text{MRS}_{cg} / \partial l}{\partial \text{MRS}_{cz} / \partial l} \frac{dg}{d\theta}.$$

Insert this in (16) to arrive at Corollary 2.

Proof of Corollary 1: When utility is given by

$$u(c, g, z, n) = \tilde{u} [w^1(c, g, z), w^2(z, n)],$$

the marginal willingness to pay for g is

$$\text{MRS}_{cg} = \frac{u'_1(\cdot) w'_g(\cdot)}{u'_1(\cdot) w'_c(\cdot)} = \frac{w'_g(c, g, z)}{w'_c(c, g, z)},$$

which is independent of n . Thus, $\partial \text{MRS}_{cg} / \partial n = 0$ and (16) reduces to the Samuelson rule. Note that $u = \tilde{u}(w(c, g), l)$ as used in Kaplow (1996) is simply a special case of the utility function above, where $w^1(c, g, z) = \tilde{w}^1(c, g)$ and $w^2(z, n) = z/n$.

B Appendix C

Special Case (Section 5): From (14)

$$\frac{dz}{d\theta} = \frac{u''_{cn} \frac{\partial T(\cdot)}{\partial \theta} - u''_{gn} \frac{dg}{d\theta}}{u''_{cn} (1 - m) + u''_{zn}}.$$

When $u = c + nw(g, z) - n \cdot h(z/n)$, $u''_{cn} = 0$. The FOC for the choice of z implies

$$h'(\cdot) = 1 - T'_z(\cdot) \implies \frac{dz}{d(1 - m)} = \frac{n}{h''(\cdot)},$$

which gives the (compensated) elasticity of earned income w.r.t. the take-home rate as

$$\varepsilon \equiv \frac{dz/z}{d(1-m)/(1-m)} = \frac{nh'(\cdot)}{zh''(\cdot)}.$$

The cross-derivative u''_{zn} then becomes

$$u_{zn} = h''(\cdot) \frac{z}{n^2} = (1-m) \frac{1}{n} \frac{1}{\varepsilon}.$$

Using this in the expression for $dR/d\theta$ (7) shows that g should be expanded as long as

$$\int u'_g(\cdot) \left(1 - \frac{m}{1-m} \varepsilon\right) f(n) dn > \text{MRT},$$

which is (17) in the main text.

References

- [1] Atkinson, A.B. and N.H. Stern (1974). “Pigou, Taxation, and Public Goods.” *Review of Economic Studies*, 41, pp. 119–28.
- [2] Auerbach, A.J. and J.R. Hines Jr. (2002). “Taxation and Economic Efficiency.” In A.J. Auerbach and M. Feldstein, *Handbook of Public Economics*, vol. 3. Amsterdam: North-Holland, pp. 1349–1421.
- [3] Ballard, C.L. and D. Fullerton (1992). “Distortionary Taxes and the Provision of Public Goods.” *Journal of Economic Perspectives*, 16, pp. 117–31.
- [4] Becker, G.S. (1965). “A Theory of the Allocation of Time.” *Economic Journal*, 75, 493–517.
- [5] Blomquist, S., V. Christiansen and L. Micheletto (2008). “Public Provision of Private Goods and Nondistortionary Marginal Tax Rates.” Mimeo.
- [6] Browning, E.K. (1987). “On the Marginal Welfare Cost of Taxation.” *American Economic Review*, 77, pp. 11–23.
- [7] Christiansen, V. (1981). “Evaluation of Public Projects under Optimal Taxation.” *Review of Economic Studies*, 48, pp. 447–57.
- [8] Dahlby, B. (1998). “Progressive taxation and the social marginal cost of public funds.” *Journal of Public Economics*, 67, 105–22.

- [9] Gahvari, F. (2006). “On the marginal cost of public funds and the optimal provision of public goods.” *Journal of Public Economics*, 90, pp. 1251–62.
- [10] Goulder, L.H., I.W.H. Parry and R.C. Williams III. (2006). “Evaluating Public Goods and Regulations: Comment.” *Journal of Economic Perspectives*, 19, pp. 233–34.
- [11] Gruber, J. and E. Saez. “The elasticity of taxable income: evidence and implications.” *Journal of Public Economics*, 84, pp.1–32.
- [12] Hylland, A. and R. Zeckhauser (1979). “Distributional Objectives Should Affect Taxes But Not Program Choice or Design.” *Scandinavian Journal of Economics*, 81, pp. 264–84.
- [13] Kaplow, L. (1996). “The Optimal Supply of Public Goods and the Distortionary Cost of Taxation.” *National Tax Journal*, 49, pp. 513–33.
- [14] Kaplow, L. (2004). “On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy.” *Journal of Economic Perspectives*, 18, pp. 159–75.
- [15] Kaplow, L. (2006). “Optimal Control of Externalities in the Presence of Income Taxation.” NBER Working Paper No. 12339.
- [16] Kaplow, L. (2008). “Optimal Policy with Heterogeneous Preferences.” NBER Working Paper No. W14170.

- [17] Kaplow, L. and S. Shavell (2002). “On the Superiority of Corrective Taxes to Quantity Regulation.” *American Law and Economics Review*, 4, pp. 1–17.
- [18] Kleven. H.J. and C.T. Kreiner (2006). “The marginal cost of public funds: Hours of work versus labor force participation.” *Journal of Public Economics*, 90, 1955–73.
- [19] Saez, E. (2001). “Using Elasticities to Derive Optimal Income Tax Rates.” *Review of Economic Studies*, 68, pp. 205–29.
- [20] Salanié, B. (2003). *The Economics of Taxation*. Cambridge: MIT Press.
- [21] Samuelson, P. (1954). “The Pure Theory of Public Expenditure.” *Review of Economics and Statistics*, 36, pp. 387–89.
- [22] Slemrod, J. and S. Yitzhaki (2001). “Integrating Expenditure and Tax Decisions: The Marginal Cost of Funds and the Marginal Benefit of Projects.” *National Tax Journal*, 54, pp. 189–201.

Chapter 4

Taxation, Regulation, and Cue-Triggered Consumption*

Nicolaj Verdelin
University of Copenhagen and EPRU

August 2007

Abstract

This paper analyzes taxation and cue regulation in the presence of cue-triggered consumption (Laibson, 2001). Exposure to a visceral cue is random and increases consumption of the cue-related good. When there are no errors in individual decision-making, it is optimal to subsidize cue-triggered consumption. When individuals are boundedly rational, it is optimal to tax sin good purchases. In any case, the optimal sin tax (or subsidy) decreases as cue regulation is tightened. This suggests that cue regulation and sin taxes are substitutable instruments. Finally, hassle policies such as smoking bans in public places are shown to be generally inefficient.

Keywords: Cue-triggered consumption; Cue regulation; Sin taxation; Bounded rationality

JEL classification: H21; D6; D87

*I am grateful to Xavier Gabaix, Andreas Kappeler, Jukka Pirttilä, as well as participants at the Nordic Workshop on Tax Policy and Public Economics 2006, and especially to Jes Winther Hansen, Henrik Jacobsen Kleven, and Claus Thustrup Kreiner for helpful comments and suggestions. All remaining errors are mine.

1 Introduction

In recent years there has been a surge in the interest of the economics profession in the role of temptation and self-control in human decision-making. This research has been motivated, as has behavioral economics in general, by the failure of traditional, neoclassical models to accurately capture central aspects of human behavior. At the same time, numerous countries have chosen to operate paternalistic policies such as sin taxes on cigarettes and advertisement bans for alcoholic beverages. However, the literature has generally failed to converge towards a unifying framework and widespread disagreement remains as to the degree of rationality on the part of human beings. Not surprisingly, different assumptions about rationality often lead to starkly contrasting policy implications.

This paper addresses a particular type of temptation-based consumption that is influenced by random external cues, or so-called *cue-triggered conditioned responses* (Laibson, 2001). In this setting, a visceral cue (the smell of freshly baked bread or the sight of somebody smoking) triggers an anticipatory response that affects demand. The temptation following cue exposure poses a challenge for welfare economics. Should it be given full hedonic value when assessing the well-being of an individual? Or is it more appropriate to assume that anticipatory emotions do not affect true long-run welfare? This comes back to the central question of the appropriate departure from full rationality. At the same time, the formal acknowledgement and inclusion of a role for external cues in determining consumption demands allows for a study of cue regulation, i.e., policies that affect the

risk of cue exposure such as legislation directed at the advertising industry. This is an important issue in practical policy about which economics has so far had little to say.

The notion of cue-triggered consumption was first formalized by Laibson (2001) in a model of addiction that follows the tradition of Becker and Murphy (1988) and assumes that choices are always aligned with true preferences. This corresponds to giving full weight to any temporary emotional response in the wake of cue exposure. In contrast, Bernheim and Rangel (2004) refer to evidence from neuroscience indicating that cues for addictive goods inhibit the proper functioning of the chemical process through which the human brain forecasts the expected pleasure from consuming. This leads to errors in individual decision-making.

The present paper provides a simplified and unified framework, encompassing both the case studied by Laibson (2001) and that considered by Bernheim and Rangel (2004), to study optimal tax policy and the role of cue regulation. This leads to new results on optimal paternalism and generates insights into the dependence of policy prescriptions on the specific modelling assumptions made.

If one assumes that individuals are fully rational it is tempting to conclude that there is no role for the government to play beyond the correction of externalities. However, the random nature of cue shocks affects the marginal utility of consumption in a stochastic sense which creates a demand for insurance. The government can supply partial insurance by subsidizing the sin good. This result extends to an environment with taste differences and also with heterogeneous earnings as long as the government has access to a non-linear

income tax. As in cases of optimal insurance design, the size of the subsidy involves a trade-off between the benefits of consumption smoothing across states and the risk of moral hazard (i.e., increased consumption of the sin good). I find that the optimal subsidy rate decreases in the degree of government cue regulation.

If instead decision utility is misaligned with true utility, standard arguments suggest using sin taxes to correct so-called *externalities*. Importantly, the present paper shows that the optimal corrective tax depends on the degree of regulation of the cue process. As regulation is tightened, the optimal sin tax decreases. Further, when targeted regulation is possible, it improves the trade-off faced by the government.

There are two main contributions of this paper. The first is to demonstrate that the nature of optimal tax policy depends crucially on the degree of individual rationality. The second is to show that cue regulation and sin taxation are substitutes, not complements, irrespective of the degree of individual rationality. Indeed, the two different treatments proposed in the literature of the short-term cravings that people experience in the wake of cue exposure lead to opposing policy recommendations. If one follows Laibson (2001) and attaches full hedonic weight to short-term cravings, the optimal policy is a *subsidy* on sin goods. If, instead, one uses the approach of Bernheim and Rangel (2004), the optimum involves a *tax* on sin goods. Naturally, the true degree of rationality is an empirical question but it seems that today there is more widespread acceptance of the notion of errors in individual decision-making. However, for the case of conditioned responses there are more complicated issues involved. Indeed, even if one accepts that demand does not

always reveal true preferences, it is not clear that conditioned responses should be given no hedonic weight. The true model may very well be a hybrid of the two polar cases.

In contrast to the model-dependent tax policy prescriptions, tighter cue regulation always leads to less reliance on the tax instrument. Tighter cue regulation decreases the risk of cue exposure and thereby reduces the overall gain from distortionary taxation: the tax (or subsidy) is aimed at helping exposed individuals at the expense of non-exposed. Thus, the interdependence between policy instruments arises because of the need to rely on distortionary taxation.

The analysis in this paper can be seen as a contribution to the literature on optimal paternalism. While still in its infancy, this literature has shown that the government may well have an important role to play in helping to correct individual errors due to bounded rationality. Most existing studies of optimal paternalistic policies have proceeded within the context of the quasi-hyperbolic discounting framework, rediscovered in the context of individual choice by Laibson (1997) and further explored by O'Donoghue and Rabin (1999, 2001). Within this model, Gruber and Köszegi (2001) derive optimal cigarette taxes while O'Donoghue and Rabin (2006) identify conditions under which optimal sin taxes increase with the average degree of irrationality in the population. Hansen (2005) presents an argument for providing public transfers in kind when agents are present-biased. Gul and Pesendorfer (2001, 2004) suggest a different interpretation of temptationist preferences based on an axiomatic, choice-theoretic approach. Their assumption of full rationality usually implies that there is no corrective role for the gov-

ernment. A notable exception is Gul and Pesendorfer (2007) who find that ‘price policies’ are generally inefficient and dominated by bans. The present paper differs from this literature by considering the role of external cues in triggering temptation. This sets the stage for an analysis of policies of cue regulation and results on the interdependencies between different policy instruments. Finally, the paper explicitly contrasts the policy prescriptions with and without individual bounded rationality.

The paper proceeds as follows. Section 2 briefly outlines the underlying idea of cue-triggered consumption while Section 3 presents the model. Section 4 explores the Laibsonian case of full rationality and presents the case for commodity specific insurance. Section 5 considers the case of errors in decision-making and analyzes optimal sin taxes while Section 6 considers cue regulation. Section 7 concludes.

2 Cue-Triggered Consumption

The starting point for this paper is the concept of cue-triggered consumption, which is the notion that external cues affect consumption decisions. A cue can be interpreted as any kind of object, whether physical or abstract, that triggers an association to consumption of a certain good. This broad interpretation naturally encompasses all kinds of direct advertising but also events such as the sight of others consuming the good or of the good itself. The idea is based on insights from cognitive neuroscience and psychology where the concepts of conditioned responses and compensatory conditioned responses have been introduced. Formally, a conditioned response is a "cue-based anticipatory response to a

physiological stimulus" and a compensatory conditioned response is a "conditioned response that is homeostatic in nature".¹ Homeostasis is an equilibrating mechanism that stabilizes the physiological system. In more mundane terms, a conditioned response is the body's preparatory reaction to a cue such as the smell of freshly baked bread. The cue triggers an anticipation of intake and induces a sense of craving. Compensatory responses arise when the body prepares itself for future intake upon cue exposure (homeostasis) as is the case with addictive drugs that have strong physiological effects. It is aimed at dampening the body's reaction to the anticipated consumption. For instance in connection with cocaine abuse, the release of dopamine, which is responsible for the 'kick', is split between the time at which anticipation of intake arises and the time of actual consumption. These insights were formalized by Laibson (2001) by introducing a random cue process and two sets of baseline preferences that are active in the states of exposure and non-exposure, respectively. Building on the original idea of Becker and Murphy (1988), Laibson (2001) introduces a cue-triggered reference-point in the basic utility function. This captures the idea that upon exposure, the marginal utility of consumption increases but the utility derived from consumption is reduced. In terms of observed behavior, this will increase the probability of consuming as well as increase the amount consumed conditional on consuming. The approach of Laibson retains the assumption of rational choice.

Bernheim and Rangel (2004) make use of a Hedonic Forecasting Mechanism (HFM)

¹These definitions are taken from a Lecture Note by David Laibson.

that responds to the brain's production of dopamine. The level of dopamine is used to forecast the pleasure from consumption, thereby helping the individual to optimize choices, and is affected by previous experiences with the good in question. This reflects the same anticipatory effect of cue exposure described above. As to decision-making, the agent may find herself in one of two states: a cold mode and a hot mode. In the cold mode, decisions are rationally governed by utility maximization. In the hot mode, the agent consumes the sin good by default. This approach is motivated by evidence that cues related to addictive goods can interfere with the proper functioning of the HFM by increasing the concentration of dopamine. This leads the agent to overestimate the pleasure from consuming the sin good.

A common trait in this literature is the focus on sin goods and negative cues, where exposure to a cue triggers a sense of craving but does not add to the experience of consuming. Indeed, Laibson considers the case when utility is adversely affected by cues. The present paper follows this tradition by restricting attention to sin goods, where consumption entails certain costs, and some of the paper's conclusions may arguably be mostly relevant for addictive goods. At the same time, I will maintain the assumption that individuals would prefer to avoid cue exposure if given the choice. However, many of the results only require that the marginal utility of consuming the sin good increases upon exposure.²

²Laibson refers the above-mentioned homeostasis when qualifying the assumption of lower utility in the state of exposure. One might also argue that since cue exposure is an external event that affects consumption choices without changing the choice set, individuals must be worse off in the state of exposure. However, keep in mind that preferences change upon exposure.

3 Model

In this section the model is laid out in its most general form allowing for income effects on consumption. This sets the stage for the analysis of the Laibsonian case. The section on bounded rationality will disregard income effects to simplify the analysis. The model is intentionally simple. Thus, intertemporal effects of consumption such as the build-up of an addictive habit are not explicitly modelled. Instead, the focus is on optimal policy and the relation between different policy instruments. The analysis in Bernheim and Rangel (2004) is an example of a full-scale model of addictive consumption. The simplified framework has the advantage of lending itself easily to several different interpretations. Therefore, many of the results would seem to have merit for many different kinds of sin goods with very different characteristics.

The model is static. The economy consists of N individuals indexed by i . Labor supply is exogenously, but individually, given and generates income y_i for agent i . Agents derive utility from consuming a composite good x and a sin good z . The budget constraint for agent i is

$$x_i + qz_i \leq y_i,$$

where q is the relative price of the sin good. The preferences of agent i are represented by the utility function

$$u_i(x_i, z_i - s \cdot r_i), \tag{1}$$

where s is an indicator variable for cue exposure, equal to 1 when exposed and 0 otherwise.

Standard assumptions are made about the shape of the utility function. Specifically, u_i is strictly concave, increasing in both its arguments, and the marginal utilities of x and z are infinite when consumption is zero.³ Cue exposure activates a reference point, r_i , in the utility function of the exposed agent, capturing the anticipatory nature of the conditioned response underlying cue-triggered consumption. This has the effect of increasing the marginal utility of consuming the sin good. In this sense cues make consuming z more tempting and r_i reflects agent i 's sensitivity to cue exposure.⁴

The process of cue exposure is random. With probability p_i agent i is subjected to a cue, which activates the individual reference point, r_i . Both the reference point and the probability of exposure is individual-specific as captured by the subscript i . This reflects that different people have different sensitivities to cues and different risks of exposure. The probability of cue exposure is likely to be partly influenced by actions under the control of the individual but this process is not modelled explicitly.

The timing of the model is as follows. First, cues are transmitted and the consumption choices are made after the resolution of uncertainty, i.e., after the state of exposure is known. The individual thus allocates expenditure across the two consumption goods

³It is perhaps more realistic, but inconsequential to the results, to allow the marginal utility of z to become negative when consumption exceeds a certain threshold to reflect that sin good consumption is associated with various costs in the form of, e.g., poorer health. A natural benchmark is then the case of separability in which $u_i(x, z) = v_i(x) + w_i(z - s \cdot r_i) - c_i(z)$, where $c_i(\cdot)$ represents the health costs of consuming the sin good. Apart from the presence of cues, this is the functional form analyzed by O'Donoghue and Rabin (2006).

⁴The form of (1) is not important for the results in this paper. Specifically, utility need not drop following cue exposure. What matters is that the marginal utility of consuming the sin good increases after exposure to a cue. However, for the purpose of discussing the merits of cue regulation in Section 7, it is assumed that cue exposure implies a loss of utility.

to maximize (1) subject to the budget constraint.⁵ This leads to individual demand functions of the form

$$x_i = x_i(q, y_i, s \cdot r_i), \quad s = 0, 1$$

$$z_i = z_i(q, y_i, s \cdot r_i), \quad s = 0, 1$$

Cue exposure leads agents to purchase more of the sin good at the expense of other commodities, i.e.,

$$z_i(q, y_i, r_i) > z_i(q, y_i, 0)$$

$$x_i(q, y_i, r_i) = y_i - qz_i(q, y_i, r_i) < y_i - qz_i(q, y_i, 0) = x_i(q, y_i, 0). \quad (2)$$

There are two competing interpretations of this set-up. Under one view it considers the agent at a moment in time and there is an instantaneous risk of exposure which will affect the consumption choice. With this interpretation the model is best understood as treating a single out of a large number of independent periods of life as in O'Donoghue and Rabin (2006). Under another view the model is a description of a long-run steady state where an agent is either exposed, which would then be the equivalent of e.g., having build an addictive habit, or the agent is not exposed. Again, in the idiosyncratic state of exposure, consumption of the sin good is higher for any agent than for that same individual in the state of non-exposure.

⁵A more general model would endogenize labor supply. Under weak conditions an agent would smooth the cue shock absorption across all consumption goods and would thus increase labor supply relative to the state of non-exposure. This is potentially important in public finance analysis but is left for future work.

Note that the model abstracts from any external effects of sin good consumption. While clearly relevant, the consequences of the presence of externalities is well understood and they are left out to focus the analysis on the novel aspects of the model. This is in line with similar work on temptation and self-control such as Bernheim and Rangel (2004) and Gruber and Köszegi (2001).

4 Welfare and the Role of Government

The government sets policy before the realization of cue shocks at a point in time when individual uncertainty as to future preferences remains. Two different types of policy instruments are considered. The government may tax (or subsidize) the sin good at the rate t , making the consumer price $q + t$, and it may promote legislation that regulates the flow of cues.⁶ Indeed, p_i is assumed to be responsive to government regulation of the cue process. This covers regulation of advertisement as well as public campaigns that affect social norms but does not include policies that restrict consumption directly such as bans. I will distinguish two types of government cue regulation: targeted and non-targeted (general). The former is aimed at a specific, identifiable group in society and affects all members of that group (although not necessarily with equal impact). The latter covers general policies that affect all members of society (although again not necessarily with equal impact). For the purpose of targeted regulation, it is assumed

⁶In this two-good economy, a subsidy on the sin good is, of course, equivalent to a tax on the composite good. This merely reflects the general point that the number of government price instruments is always no larger than the number of commodities minus one. More importantly, with endogenous labor supply it may, under certain conditions, be possible to implement the optimum relying only on an income tax.

that the government observes a number of individual traits, e.g., age and gender, and is able to target regulatory policies at groups of individuals who share a certain trait. Groups are indexed from 1 to M and agent i is said to belong to group J if $i \in J$. Regulation is assumed to be measurable on a well-defined scale, and γ_g denotes the degree of general regulation, whereas γ_J denotes the degree of regulation targeted at group J . The probability of individual cue exposure is then a function of the degree of general regulation and the degree of targeted regulation that affects the groups to which the individual belongs:

$$p_i = p_i(\gamma_g, \gamma_1, \dots, \gamma_M)$$

with $\partial p_i / \partial \gamma_g < 0$, $\partial p_i / \partial \gamma_J < 0$ if $i \in J$ and $\partial p_i / \partial \gamma_J = 0$ if $i \notin J$. Note that tighter cue regulation translates into a lower risk of cue exposure.

The model laid out in section 3 involves each agent having two different utility functions depending on the idiosyncratic state of mind. This raises the issue of what is an appropriate welfare criterion. Two competing assumptions have been suggested in the literature. In the model of Laibson (2001), full weight is given to any sense of craving and the agent is assumed to be fully rational as in Becker and Murphy (1988). In this case, utility is truly state-dependent and there is never any conflict between choices and self-interest. Thus, the appropriate welfare criterion takes into account the reference point when the agent has been subjected to a cue. Section 5 deals with this situation of full rationality.

In the analysis of Bernheim and Rangel (2004), cue exposure is assumed to drive the

agent to act irrationally. Thus, while an exposed individual behaves as if preferences were state-dependent, true utility is in fact represented by (1) with $s = 0$. This causes individual choices to diverge from the self-interest of the agent whenever she has been exposed to a cue. Under this view the government should apply a welfare criterion that is independent of the idiosyncratic state of mind of the individual agent.⁷ This leads to a case for paternalism which is the subject of Section 6.

It is assumed throughout the paper that the government maximizes a utilitarian social welfare function, and I will use g_i^e and g_i to denote the marginal social welfare weights attached to individual i in the exposed and non-exposed state, respectively.

In the present model cues are always detrimental to individual welfare. A more complete description would address the issue of the benefits of cues as well as the direct costs of implementing and enforcing cue regulation. However, since these questions are as yet not well understood in the literature, the first part of this paper will focus only on the relation between cue regulation and tax policy. Section 7 contains a less formal treatment of the question of optimal cue regulation.

5 The Case of Full Rationality

In this section it is assumed that decision-making is always aligned with true preferences. Further, unless otherwise stated, it is assumed that all agents have the same income, $y_i = y$ for all i . Consequently, there is no corrective role for the government and it

⁷See the chapter by Bernheim and Rangel (2007) on welfare evaluations with bounded rationality and a discussion specifically within the context of their 2004 model.

is tempting to conclude that the optimal policy has $t = 0$. However, cue exposure is a random event which affects the marginal utility of income. Indeed, it follows from (2), using the Envelope Theorem, the assumption of full rationality and concavity of the individual utility function, that

$$\frac{\partial V_i}{\partial y_i} = u_{ix}(x_i(q, y_i, s \cdot r_i)) \Rightarrow \frac{\partial V_i}{\partial y_i} \Big|_{s=1} > \frac{\partial V_i}{\partial y_i} \Big|_{s=0}, \quad (3)$$

where $V_i(q, y_i, s \cdot r_i)$ is the indirect utility function. Naturally, a government that respects individual preferences will set $g_i^e > g_i$ for all i . The marginal utility of income is higher in the state of exposure because of the temptation to consume z , which leads the consumer to divert resources away from spending on the composite good. The divergence of marginal utilities of income across states creates a desire, ex ante, to transfer resources from the state of non-exposure to that of exposure. This is similar to insurance against an adverse shock. In the case of cigarettes, alcohol or drug abuse or gambling, the stakes can be quite high. Conceivably, no private insurance company will provide this kind of insurance due to obvious problems of adverse selection under asymmetric information about the state of exposure.⁸ This leaves room for the government to beneficially intervene and provide insurance against "affective consumption", a point first made by Bernheim and Rangel (2004).

⁸A caveat is in order. If the agent has access to perfect capital markets, she may self-insure by borrowing against future earnings. This is not possible in the static framework presented here, but when the model is interpreted as representing a single out of a large number of periods, intertemporal consumption smoothing is possible. Even in this case, if individual uncertainty remains, agents will still be better off by pooling idiosyncratic risks. These benefits then have to be weighted against the distortionary effects of taxation. However, under the preferred steady state interpretation of the model, where the state of exposure is in fact a permanent state of addiction, this caveat has no bearing. I am grateful to Claus Thustrup Kreiner for pointing this out.

The case for government intervention does not rest on an assumed informational advantage over the private insurance market. Even if the government is unable to verify whether an individual has been exposed to a cue, it may exploit the difference in consumption patterns between exposed and non-exposed individuals to effectively provide insurance. The crucial observation is that sin good consumption is higher for exposed individuals. It follows that a policy of subsidizing sin good consumption will provide commodity specific insurance against the risk of cue exposure.

It is instructive to begin by considering the case of a single (ex ante) representative agent. The government sets the tax rate to maximize expected utility subject to a budget constraint

$$\begin{aligned} \max_t & pV(t, y, s \cdot r) + (1 - p)V(t, y, 0) \\ \text{s.t. } & t[pz^e(t) + (1 - p_i)z(t)] = \pi, \end{aligned}$$

where superscript ‘e’ denotes exposed, and π is a lump sum tax. It can be shown that, starting from $t = 0$, a small subsidy improves the welfare of the representative agent whenever

$$[z^e(0) - z(0)] \left(\frac{\partial V^e}{\partial y} - \frac{\partial V}{\partial y} \right) > 0,$$

which we know from above is always satisfied, since both consumption and the marginal utility of income is higher in the state of exposure. A subsidy effectively transfers income from non-exposed to exposed individuals because the value of the subsidy is higher for this group than the value of the lump sum tax. At the same time, a small subsidy

introduces only a negligible distortionary cost when prices are initially undistorted.

When there is heterogeneity in preferences, matters are more complicated. While the individual demand for insurance continues unabated, there is no longer any guarantee that a subsidy strictly targets the intended group of exposees. Essentially, a subsidy to the sin good no longer serves as an actuarially fair insurance contract for everyone. The expected (monetary) gain to individual i from a small subsidy is

$$p_i z_i^e + (1 - p_i) z_i - \bar{z},$$

where \bar{z} denotes average consumption of the sin good in the population and is the equivalent of the insurance premium. With homogeneous preferences this expression is zero and the implicit insurance contract is actuarially fair. When preferences are diverse, subsidizing the sin good only provides insurance to those individuals whose consumption of the sin good is higher than average. Further, the attractiveness of the implicit insurance contract depends on the individual risk of cue exposure. Thus, a subsidy is more beneficial to individuals with a high consumption of the sin good and with a high risk of cue exposure.⁹ Indeed, the value of the implicit insurance to individual i is

$$p_i \frac{\partial V_i^e}{\partial y_i} (z_i^e - \bar{z}) + (1 - p_i) \frac{\partial V_i}{\partial y_i} (z_i - \bar{z}).$$

⁹The analysis so far has focused on the case of continuous consumption. It may well be argued that for many sin goods the extensive margin is more important than the intensive margin. While there is arguably some dispersion in daily cigarette consumption among smokers, the more important margin appears to be the decision of whether to smoke or not. The results in this section are robust to an alternative model where the extensive margin is the only margin of response. In this case, the government's trade-off will be better if the distribution of gains from consumption is more dispersed. This opens for the possibility of subsidies that in certain ranges have little or no effect on consumption decisions.

In spite of these complications, there is still a strong case for subsidizing purchases of z as, on average, consumption of the sin good is higher in the state of exposure. If we make the reasonable assumption that the government has no intrinsic preference for individual sin good consumption, it is always optimal to subsidize cue-triggered consumption whenever the following condition holds

$$\sum_{i \in N} [p_i(\gamma_g, \gamma_1, \dots, \gamma_M)(g_i^e - g_i)(z_i^e - \bar{z}_I)] > 0, \quad (4)$$

for $g_i^e > g_i$. This condition ensures that (a) the government has no intrinsic preference for sin good consumption among exposees that overrules the individual demand for insurance and (b) the risk of cue exposure is positively correlated with sin good demand in the exposed state.¹⁰

Proposition 1. *A government with no intrinsic preference for individual sin good consumption will always set $t^* < 0$ if (4) is satisfied for $g_i^e > g_i$.*

Further, the optimal subsidy rate is decreasing in the degree of cue regulation, whether general (γ_g) or targeted at group J (γ_J), whenever preferences satisfy a single-crossing condition.

Proof. See Appendix A.

An important question concerns the interdependence between the different government instruments. Proposition 1 establishes that the optimal subsidy rate decreases (i.e.,

¹⁰Part (b) merely requires that cue exposure has a sufficiently strong impact on sin good purchases to make the demand for insurance high enough.

the tax increases) with the degree of cue regulation. Intuitively, when regulation is increased the frequency of cue exposure decreases such that the risk of experiencing a high marginal utility of income is reduced. The fact that fewer people become exposed implies that the gain from subsidizing the sin good is reduced. As a consequence, the optimal subsidy rate declines as regulation is intensified. In this sense, cue regulation and taxation are substitutable instruments. As always, comparative statics under heterogeneous preferences are complicated by the need to ensure global optimality when social welfare is not necessarily strictly concave in the tax rate. This explains the necessary single-crossing condition (A-1), which basically assumes that for any t the group of exposees is always more hurt at the margin by higher taxation than the group of non-exposees.¹¹

I conclude this section by briefly considering the case of income heterogeneity. As is well-known from the literature on optimal commodity taxation (see, e.g., Saez, 2002), it is generally unnecessary to rely on commodity taxes when the government has access to a non-linear income tax. The case for commodity specific insurance presented in this paper is not based on a desire to redistribute across income groups, but rather between idiosyncratic states of nature. Therefore, it is possible to show that the argument for a subsidy to z is entirely robust to the inclusion of income heterogeneity as long as an optimal non-linear income tax is in place. Importantly, this implies that even if low

¹¹This is further complicated by the presence of income effects. Essentially, (A-1) also assumes that if a sum of money were to be spread evenly across the population, the optimal subsidy would not increase too much. Increased regulation improves the net fiscal status of both exposed and non-exposed individuals. This may change the relative marginal utilities of income of the two groups. In addition, the lower frequency of exposure makes it less costly to transfer resources to the exposed state which also tends to increase the subsidy. Finally, the marginal distortionary costs may change as consumption changes due to income effects and the higher fraction of non-exposees.

income individuals suffer a greater risk of cue exposure or are more sensitive to cues, the optimal sin subsidy is not affected. Redistribution between income groups should be entirely handled through the income tax. In contrast, redistribution across states is only possible by use of a commodity subsidy.

6 The Case of Bounded Rationality

In this section, the change in behavior upon cue exposure is attributed to bounded rationality as in Bernheim and Rangel (2004). Thus, the short term feeling of temptation that results from exposure to a cue is considered an ‘illusion of the brain’. In chemical terms, the brain’s release of dopamine as a predictor of the pleasure of intake is disturbed by the presence of the cue. This causes the brain to overestimate the pleasure from consuming. As a consequence, if left to herself, an agent’s choices will be suboptimal in the exposed state of mind. Specifically, the agent acts as if true utility were given by (1) and thus were state-dependent. But in fact, utility does not vary with the state.¹² Thus, independently of cue exposure, welfare is given by $u_i(x, z)$, corresponding to $s = 0$, while decisions are made in order to maximize (1) after the realization of the idiosyncratic state. It is clear that this leads the agent to overconsume the sin good in the exposed state. This creates scope for government intervention. In the previous section, this was motivated by a demand for insurance that could not conceivably be satisfied by the private market.

The case for insurance continues to hold as cue exposure adversely affects the disposable

¹²This can be considered a normalization. It is straightforward to allow for some state-dependence of true preferences but assume that the perceived reference point is higher than the true reference point.

income for goods other than the sin good. But at the same time, there is a need to correct an error in individual choices. With perfect information this trade-off would not exist but when the government cannot condition commodity taxes and transfers on exposure, the concern for income distribution would likely affect the results. However, the subsequent analysis makes the simplifying assumption that utility is quasi-linear, eliminating the demand for insurance.¹³

The utility of agent i is given by

$$u_i(x, z) = w_i(z) + x$$

Meanwhile, consumption choices are made to maximize

$$w_i(z - s \cdot r_i) + x,$$

which deviates from true utility whenever the agent has been subjected to a cue i.e., when $s = 1$. When $s = 0$ there is no conflict between decision utility and true utility. It is assumed that the first-order conditions yield an interior solution for all i and in both states.

The government faces a problem of asymmetric information as cue exposure is unobservable and preferences are heterogeneous. Thus, a high consumption of the sin good can be due to either cue exposure or an innately strong preference for the good. The government is, therefore, left with the option of affecting behavior through the use of a

¹³I allow for income heterogeneity in this section, but the assumptions of quasi-linearity and utilitarian social welfare imply that income distribution is of no importance. Of course, with a concave social welfare function the government would still wish to redistribute to those with low earnings.

distortionary commodity tax or cue regulation. As first shown by O’Donoghue and Rabin (2001) in the context of quasi-hyperbolic discounting, it is optimal for the government to apply at least a small positive tax rate on sin good consumption as long as the demand of exposed individuals is not entirely inelastic. This result easily extends to the present model. Proposition 2 presents the case for an optimal sin tax along with results on the dependence of the optimal tax rate on the degree of regulation.

Proposition 2. *Assume utility is quasi-linear and all individuals choose an interior solution. Then, as long as w_i is strictly concave for all i , the optimal sin tax rate is*

- (a) *Positive, $t^* > 0$.*
- (b) *Decreasing in γ_g .*
- (c) *Decreasing in γ_J for all J .*

Proof. See Appendix B.

Proposition 2 shows that it is optimal to use a distortionary commodity tax to at least partially correct the ‘internality’ caused by cue exposure. This is an example of what the literature has coined ‘asymmetric paternalism’ (O’Donoghue and Rabin 2003).¹⁴ At a zero tax rate, the true average (for each type) marginal rate of substitution between the sin good and all other goods is less than the price because of the overconsumption of exposees. By raising the tax rate slightly, no harm is done to non-exposees (to a first order) but the welfare of exposees is strictly increased as long as consumption is not

¹⁴Other terms have been put forth. Thaler and Sunstein (2003) propose ‘Libertarian Paternalism’ while Choi et al. (2003) suggest ‘Benign Paternalism’.

entirely price inelastic. In any case, taxation interferes with the choices of non-exposed which is at the heart of the trade-off faced by the government.

A second lesson from Proposition 2 is that there is a close relation between regulation and the optimal sin tax. If the government were to increase regulation of cue emission at a general level, the optimal sin tax rate would decrease. In this sense, sin taxes and cue regulation are substitutes, not complements. The reason is simply that cue regulation ameliorates the problem of overconsumption by lowering the incidence of cue exposure. Thus, the marginal gain from taxation is reduced. In reality, it seems to be more frequently the case that countries with stricter regulation also implement high tax rates. According to the analysis above, this policy is ill-guided. Of course, this is only casual observation and any observed correlation could, for instance, be due to low regulation, low taxation countries having implemented insufficient measures. Part (c) of the result shows that regulatory policies aimed only at specific groups should still affect the general sin tax. This points to a crucial advantage of targeted (and non-targeted) regulation: it improves the trade-off faced by the government by easing the need to rely on the indiscriminate tax instrument. Taxation has the benefit of only affecting the market price, leaving consumption choices to individual agents, but at the same time forces everyone to face the same tax rate regardless of differences in preferences for the sin good, thus intervening too much from the viewpoint of certain groups and too little from the viewpoint of others. However, to assess the overall desirability of using cue regulation, whether targeted or not, requires knowledge of the potential costs of cue

regulation. Similarly, it is crucial to consider the feasibility of affecting cue exposure through regulatory policies, in particular if these are aimed only at certain groups.

7 Cue Regulation

One of the most important advantages of the theory of cue-triggered consumption is that it provides a framework for discussing policies such as regulation of the advertising industry. A standard theory of consumer choice would have nothing meaningful to say about the merits of cue regulation. However, although providing a theory of the effects of external cues on choice behavior is a crucial step toward a theory of optimal cue regulation, the most essential element remains the identification of the welfare effects of cue exposure. It is beyond the scope of this paper to develop a theory of the cost or benefits of cues in terms of, e.g., the conveying of information or the nuisance of interrupted TV-shows. Rather, the analysis in the present section builds on the formalized notion that cue exposure, by altering individual choices, adversely affects the well-being of exposees. The analysis centers around the results on the interdependency between taxation and cue regulation presented in Sections 5 and 6. When combined with a discussion of the process of cue regulation and the potential costs and benefits of cues, this sets the stage for a first attempt at identifying the optimal level of cue regulation. However, because of the speculative and informal nature of this discussion, the present section constitutes by no means a complete account of optimal cue regulation. Meanwhile, it remains an important task for future research to identify the costs and benefits of cues,

e.g., through observations of individual cue management.

The most apparent advantage of cue regulation over sin taxation is that it does not distort prices. Thus, non-exposed individuals are unaffected by the policy as long as cues do not serve a role as valuable providers of information. But in contrast to the tax instrument, cue regulation neither provides insurance nor corrects decision-errors for those who become exposed. It follows that an important consequence of the optimal tax-regulation interdependency is that exposed individuals are made worse off, *ex post*, when regulation is tightened. Therefore, unless cue regulation is able to eliminate temptation-driven consumption binges altogether, it can never entirely replace the sin tax. The relationship between regulation and taxation mainly arises because of the need to rely on a distortionary tax instrument to correct decisions or provide insurance, not because cue regulation is a direct substitute for corrective taxation. Indeed, stricter cue regulation only affects the optimal sin tax when consumption in the non-exposed state is price-sensitive or if some individuals are more sensitive to general regulation than others. In fact, the smaller is the demand response of non-exposed, the smaller is the impact of cue regulation on optimal sin taxation. This relationship also holds when individuals are fully rational, but in a weaker form, since the subsidy also distorts the consumption of exposees. In terms of practical policy, the tax-regulation interdependency implies, for instance, that the widespread use of advertising bans and designated smoking areas should be accompanied by a decrease in excise taxes on cigarettes (whenever these are initially optimal).

In a world of heterogeneous people, cue patterns are bound to be very diverse, that is different individuals are subjected to different types of cues. In many cases, cue patterns follow a number of observable traits. For instance, youngsters are highly exposed to TV-commercials and less so to advertising in newspapers, which, in contrast, provides an effective forum for reaching certain parts of the adult population. The government may take advantage of the difference in cue patterns by targeting cue regulation at certain groups in society. For instance, preventing producers of alcoholic beverages from engaging in sports sponsorships reduces the risk of cue exposure for the sports-interested part of the population. Similarly, a ban on advertisement for unhealthy foods and candy in commercial blocks surrounding TV programs aimed at children provides cue management for this group. By reducing reliance on the indiscriminate tax instrument, targeted regulation can help solve a great inefficiency because the tax must be set to strike a balance between the interests of a heterogeneous population.¹⁵ Thus, group-specific regulation will be most beneficial when aimed at a group whose preferences for the sin good differ substantially from those of the rest of the population. An example where targeted regulation is both feasible and likely to entail large gains is the case of young people and a wide range of sin goods such as alcohol and cigarettes. Stretching the interpretation of the model somewhat, the indirect utility function may be seen as the expected present discounted value of the future stream of consumption choices with today's choice of sin

¹⁵This trade-off would be improved if the government had access to group-specific taxation. However, in practice group-differentiated commodity taxes are almost never observed - most likely because of obvious enforcement problems.

good consumption entailing long run effects. For young people, the present and future values from consuming the sin good are highly uncertain, which translates into a utility function that is very concave in the sin good. This makes excessive consumption more damaging and calls for highly intrusive corrective measures aimed at this group. But unlike the tax instrument, which invokes a cost on everyone else, targeted cue regulation affects only the intended group.¹⁶

As with any type of regulation there will be attempts at circumventing the legislation whenever advertising for a product is prohibited. This constitutes an important restriction on the effectiveness of regulation that should be taken into account when assessing the merits of a policy proposal. It is likely that enforcement may be a bigger problem for targeted regulation, limiting the scope for using this type of legislation

The discussion so far has focused on cue reduction policies such as advertising bans. However, policymaking in recent years has witnessed an increasing reliance on what could be labelled ‘hassle’ policies. An example of a hassle policy is the banning of smoking in public places. What defines hassle policies is not so much their ability to reduce the risk of cue exposure (which may be limited) but rather that they reduce the utility from consuming the sin good. It is possible to address the desirability of hassle policies within the current model by making a small extension. Assume, similar to the case of cue

¹⁶The example of children and young adults points to situations when the optimal policy may extend as far as banning the sale of specific sin goods to certain groups. A ban is only relevant when the costs of excessive consumption become so large that they exceed the benefits of rational or non-exposed intake. This scenario is most likely when the consumption of non-exposed is very limited. Obvious examples include various types of narcotics.

reduction policies, that the government has access to a set of hassle policies, which may be targeted at specific groups or directed at the entire population. I will use β to denote a hassle policy such that β_g is general hassles whereas β_J covers hassle policies aimed at group J (a higher value of β implies more hassles). The hassle effect for individual i is captured by defining a hassle function, $h_i = h(\beta_g, \beta_1, \dots, \beta_M) \in (0, 1]$ with $h_{\beta_g}, h_{\beta_J} < 0$, that is multiplied on sin good utility. Thus, state-dependent utility becomes

$$u_i [x_i, h(\beta_g, \beta_1, \dots, \beta_M) \cdot v_i(z_i - s \cdot r_i)], \quad s = 0, 1, \quad (5)$$

where $v_i(a) = a$ except when utility is quasi-linear, in which case $v_i(a) = w_i(a)$. Apart from this simple extension to the utility function, the model is as before. In particular, I will continue to distinguish two cases based on the assumed degree of individual rationality. It is clear from (5) that hassle policies lower sin good demand as well as utility for the affected groups. Are hassle policies then ever a part of the optimum? Proposition 3 provides the answer.¹⁷

Proposition 3. *General hassle policies are always dominated by sin taxation. Targeted hassle policies can only ever be optimal when individuals are boundedly rational, and then only for groups where $t_J^* \gg t^*$ with t_J^* being the optimal group-specific tax for group J .*

Proof. See Appendix C.

The intuition for this result is straightforward. The distortion imposed by a sin tax is entirely captured by its effect on compensated sin good demand. If sin good consumption

¹⁷A similar result in the context of illegal drugs appears in Becker et al. (2006).

is initially too high, a sin tax may improve welfare. The utility loss that results from individual consumers having to pay a higher price on all inframarginal units of consumption is only partial, as the proceeds from the tax are at the government's disposal. Therefore, income effects have no welfare consequences. In contrast, hassle policies induce a utility loss on individual consumers that is not compensated by income transfers. As a result, regardless of the degree of individual rationality and for any functional forms for u_i , sin taxation always dominates general hassle policies.

The only feature giving hassle policies a potential advantage over sin taxation is the ability to direct the policy at specific groups. Under full rationality any policy that corrects consumption choices is undesirable. When individuals are boundedly rational to varying degrees, the unconstrained optimum would involve selective taxation. However, as this is ruled out by assumption, it is possible that a second-best policy of directed hassles can raise overall welfare. But only if the targeted policy affects a group of individuals for whom the optimal sin tax provides too little correction, and then only if the problem of overconsumption is sufficiently severe.

There are a few caveats to this analysis. First, hassle policies may provide an additional means to lower the risk of cue exposure. Second, they may be effective in overcoming externalities such as second-hand smoke, although, in principle, internalities are not very different from externalities such that similar arguments may apply.

8 Conclusion

The present paper has explored optimal taxes on cue-triggered consumption goods under two different assumptions about the degree of rationality in individual decision-making. The policy recommendations from the fully rational Laibsonian case are in stark contrast to those derived within a framework of bounded rationality. In particular, sin goods should be subsidized under fairly general conditions when there are no errors in decision-making. In contrast, when cue exposure increases sin good purchases above the individually optimal amount, a corrective sin tax is optimal. In either case, reliance on the price instrument should be reduced as cue regulation is tightened. This implies, for instance, that stricter laws regulating advertising by cigarette companies should be accompanied by lower excise taxes on cigarettes, assuming the tax is initially at its optimal level. Furthermore, hassle policies, such as smoking bans in public places, are generally dominated by sin taxation. This indicates that the widespread use of hassles in recent years may be ill-advised unless there are political constraints limiting the scope for using excise taxes. Of course, there may be real-world phenomena not captured by the very stylized model presented in this paper. It may be, for instance, that the population at large fails to understand the full damaging effects of consuming the sin good and that the two signals, regulation and taxation, reinforce each other in the attempt to inform the public of the true costs of consumption. Further work on the optimal use of different forms of regulation is clearly needed.

Appendix A

Proof of Proposition 1. The government objective function is

$$W = \sum_{i \in N} \left[p_i (\gamma_g, \gamma_1, \dots, \gamma_M) \alpha_i^e V_i^e (t, \pi) + [1 - p_i (\gamma_g, \gamma_1, \dots, \gamma_M)] \alpha_i V_i (t, \pi) \right],$$

where α_i^e and α_i denotes the social welfare weights for exposed and non-exposed type i 's, respectively. The government budget constraint is

$$\pi + t\bar{z} (t, \pi) = 0,$$

where π is a lump sum tax and $\bar{z} \equiv \frac{1}{|N|} \sum_{i \in N} [p_i z_i^e + (1 - p_i) z_i]$ is average demand for the sin good.

The total derivative of W w.r.t. t , taking into account the budget effect on π , is

$$\begin{aligned} \frac{dW}{dt} &= \sum_{i \in N} \left[p_i (\gamma_g, \gamma_1, \dots, \gamma_M) \alpha_i^e \frac{dV_i^e (t, \pi)}{dt} + [1 - p_i (\gamma_g, \gamma_1, \dots, \gamma_M)] \alpha_i \frac{dV_i (t, \pi)}{dt} \right] \\ &= \sum_{i \in N} \left[\begin{array}{l} -p_i (\gamma_g, \gamma_1, \dots, \gamma_M) \alpha_i^e \frac{\partial V_i^e (t, \pi)}{\partial y_i} (z_i^e + \frac{d\pi}{dt}) \\ - [1 - p_i (\gamma_g, \gamma_1, \dots, \gamma_M)] \alpha_i \frac{\partial V_i (t, \pi)}{\partial y_i} (z_i + \frac{d\pi}{dt}) \end{array} \right] \end{aligned}$$

where use is made of Roy's identity and $d\pi/dt = -(\bar{z} + t \frac{\partial \bar{z}}{\partial t}) / (1 + t \frac{\partial \bar{z}}{\partial \pi})$.

Denote by $g_i^e \equiv \frac{1}{\lambda} \alpha_i^e \frac{\partial V_i^e}{\partial y_i}$ and $g_i \equiv \frac{1}{\lambda} \alpha_i \frac{\partial V_i}{\partial y_i}$ the social marginal utilities of income for exposed and non-exposed type i 's, respectively. Assume that the government has no intrinsic preference for individual consumption of the sin good (Saez, 2002), i.e., g_i is such that whenever $g_i^e = g_i$

$$\sum_{i \in N} p_i (\gamma_g, \gamma_1, \dots, \gamma_M) g_i (z_i^e - \bar{z}) + \sum_{i \in N} (1 - p_i (\gamma_g, \gamma_1, \dots, \gamma_M)) g_i (z_i - \bar{z}) = 0,$$

and the second-order conditions are met such that $t = 0$ is optimal when $g_i^e = g_i$. Then it follows that for $t = 0$

$$\frac{1}{\lambda} \frac{dW}{dt} = - \sum_{i \in N} p_i(\gamma_g, \gamma_1, \dots, \gamma_M) g_i^e (z_i^e - \bar{z}) - \sum_{i \in N} (1 - p_i(\gamma_g, \gamma_1, \dots, \gamma_M)) g_i (z_i - \bar{z}) < 0,$$

whenever $g_i^e > g_i$ and

$$\sum_{i \in N} p_i(\gamma_g, \gamma_1, \dots, \gamma_M) (g_i^e - g_i) (z_i^e - \bar{z}) > 0.$$

The latter condition ensures that a) on average there is a demand for commodity specific insurance and b) the government has no intrinsic preference for sin good consumption among exposees. It follows that a government with no intrinsic preference for individual sin good consumption will optimally set $t^* < 0$ whenever $g_i^e > g_i$ for all i .

Now consider the interdependency between general regulation and the optimal sin tax rate. The derivative, holding t constant, of dW/dt w.r.t. γ_g is

$$\begin{aligned} \frac{d(dW/dt)}{d\gamma_g} &= \sum_{i \in N} \frac{\partial p_i(\gamma_g, \gamma_1, \dots, \gamma_M)}{\partial \gamma_g} \left[\alpha_i^e \frac{dV_i^e(t, \pi)}{dt} - \alpha_i \frac{dV_i(t, \pi)}{dt} \right] \\ &\quad + \sum_{i \in N} \frac{d\pi}{d\gamma_g} \left[p_i(\gamma_g, \gamma_1, \dots, \gamma_M) \alpha_i^e \frac{d(dV_i^e(t, \pi)/dt)}{d\pi} + [1 - p_i(\gamma_g, \gamma_1, \dots, \gamma_M)] \alpha_i \frac{d(dV_i(t, \pi)/dt)}{d\pi} \right]. \end{aligned}$$

The first summation captures the direct effect of increasing γ_g (disregarding income effects). It reflects that tighter cue regulation lowers the risk of exposure and changes the composition between exposees and non-exposees. A sufficient condition for t^* to be increasing in γ_g is that $d(dW/dt)/d\gamma_g > 0$ for all t . When this is the case, an increase in γ_g must imply that there exist a $t > t^*$ that now dominates t^* , and no $t' < t^*$ can dominate t^* if t^* was originally optimal. The latter follows since dW/dt is now higher for

all t . It follows that

$$\begin{aligned} & \sum_{i \in N} \frac{\partial p_i(\gamma_g, \gamma_1, \dots, \gamma_M)}{\partial \gamma_g} \left[\alpha_i^e \frac{dV_i^e(t, \pi)}{dt} - \alpha_i \frac{dV_i(t, \pi)}{dt} \right] \\ & > -\frac{d\pi}{d\gamma_g} \sum_{i \in N} \left[p_i(\gamma_g, \gamma_1, \dots, \gamma_M) \alpha_i^e \frac{d(dV_i^e(t, \pi)/dt)}{d\pi} + [1 - p_i(\gamma_g, \gamma_1, \dots, \gamma_M)] \alpha_i \frac{d(dV_i(t, \pi)/dt)}{d\pi} \right] \end{aligned} \quad (\text{A-1})$$

for all t is a sufficient single-crossing condition to ensure that the optimal subsidy rate decreases with γ_g . A similar argument can easily be made for γ_J for any J . Apart from the complexity introduced by the presence of income effects, (A-1) basically amounts to assuming that for any t the group of exposees is always more hurt at the margin by taxation than the group of non-exposees (the left-hand side of (A-1) is positive). Indeed, this assumption is sufficient to satisfy (A-1) if the marginal welfare effect of taxation increases with income (corresponding to the right-hand side of (A-1) being negative as $d\pi/d\gamma_g < 0$). Q.E.D.

Appendix B

Proof of Proposition 2. (a) The government solves

$$\begin{aligned} \max_{t, \{\pi_i\}_{i \in N}} W &= \sum_{i \in N} \{p_i(\gamma_g, \gamma_1, \dots, \gamma_M) V_i^e(t, \pi_i) + [1 - p_i(\gamma_g, \gamma_1, \dots, \gamma_M)] V_i(t, \pi_i)\} \\ \text{s.t.} \quad & \frac{1}{|N|} \sum_{i \in N} \pi_i + t\bar{z}(t, \pi) = 0 \\ & \pi_i = \pi_j \quad \forall i, j \text{ with } y_i = y_j, \end{aligned}$$

where π_i is an income specific lump-sum tax. By the quasi-linearity of individual utility and the utilitarian shape of the social welfare function, income distribution is of no

importance as long as everyone chooses an interior solution. Indeed, using the balanced budget condition (the arguments in $p_i(\cdot)$ have been suppressed to simplify notation), one can show that

$$\begin{aligned} W &= \sum_{i \in N} \{p_i V_i^e(t, \pi_i) + (1 - p_i) V_i(t, \pi_i)\} \\ &= \sum_{i \in N} \{p_i (w_i(z_i^e(t)) + y_i - qz_i^e(t)) + (1 - p_i) [w_i(z_i(t)) + y_i - qz_i(t)]\}. \end{aligned}$$

The derivative of W w.r.t. t is

$$\begin{aligned} \frac{dW}{dt} &= \sum_{i \in N} \left[p_i (w_{iz}(z_i^e(t)) - q) \frac{dz_i^e(t)}{dt} + (1 - p_i) (w_{iz}(z_i(t)) - q) \frac{dz_i(t)}{dt} \right] \\ &= \sum_{i \in N} [p_i w_{iz}(z_i^e(t)) + (1 - p_i) w_{iz}(z_i(t)) - q] \frac{dz_i^e(t)}{dt} \end{aligned}$$

as $\frac{dz_i^e(t)}{dt} = \frac{dz_i(t)}{dt}$ from the relation $z_i^e = z_i + r_i$. Evaluated at $t = 0$ this becomes

$$\begin{aligned} &\sum_{i \in N} \left[p_i (w_{iz}(z_i^e(0)) - q) \frac{dz_i^e(0)}{dt} + (1 - p_i) (w_{iz}(z_i(0)) - q) z_i(0) \frac{dz_i(0)}{dt} \right] \\ &= \sum_{i \in N} p_i (w_{iz}(z_i^e(0)) - q) \frac{dz_i^e(0)}{dt} \\ &> 0, \end{aligned}$$

as $\frac{dz_i^e(0)}{dt} < 0$ by the strict concavity of w_i and the assumption of an interior solution for all i . From the individual FOC, $w_{iz}(z_i^e(0) - r_i) - q = 0 \implies w_{iz}(z_i^e(0)) - q < 0$. It follows by continuity, and the fact that $dW/dt > 0$ for all $t < 0$, that the optimal tax rate is positive.

(b) Taking the derivative of dW/dt w.r.t. γ_g , holding t constant, gives

$$\frac{d(dW/dt)}{d\gamma_g} = \sum_{i \in N} \frac{\partial p_i(\gamma_g, \gamma_1, \dots, \gamma_M)}{\partial \gamma_g} \{w_{iz}[z_i^e(t)] - w_{iz}[z_i(t)]\} \frac{dz_i^e(t)}{dt} < 0,$$

for all t by the strict concavity of w_i and the fact that $z_i^e(t) > z_i(t)$ for all t . It follows that the optimal sin tax rate must decrease with γ_g since some $t < t^*$ now dominates t^* , and no $t' > t^*$ can dominate t^* if t^* was originally optimal. The latter follows since dW/dt is now lower for all t .

(c) The same procedure as under (b) shows that

$$\frac{d(dW/dt)}{d\gamma_j} = \sum_{i \in J} \frac{\partial p_i(\gamma_g, \gamma_1, \dots, \gamma_M)}{\partial \gamma_j} \{w_{iz}(z_i^e(t)) - w_{iz}[z_i(t)]\} \frac{dz_i^e(t)}{dt} < 0.$$

Q.E.D.

Appendix C

Proof of Proposition 3. General hassle policies function exactly like a tax in terms of consumption responses. Thus, hassles cause sin good demands to fall. But unlike the tax, hassle policies induce a utility loss on all inframarginal units of consumption. Thus, general hassle policies are always dominated by a sin tax.

Targeted hassle policies can never be optimal under full rationality, as choices are optimized and hassles induce a utility loss. In the case of bounded rationality, the effect on social welfare of a hassle directed at group J , $d\beta_J$, is

$$dW = \left\{ \underbrace{\sum_{i \in J} h_{\beta_J} [p_i w_i^e + (1 - p_i) w_i]}_{< 0} + \sum_{i \in J} [p_i (h w_{iz}^e - q) dz_i^e + (1 - p_i) (h w_{iz} - q) dz_i] \right\} d\beta_J.$$

It follows that $dW/d\beta_J > 0$ requires that the last summation is positive. If the optimal directed tax for group J is *smaller* than the optimal general sin tax, $t_J^* < t^*$, it must be

the case that, when evaluated at $t = t^*$ (and $t_J = 0$), a small tax only on group J lowers welfare, i.e.,

$$dW = \sum_{i \in J} [p_i (hw_{iz}^e - q) dz_i^e + (1 - p_i) (hw_{iz} - q) dz_i] dt_J < 0,$$

such that directed hassles require $t_J^* > t^*$. Further, t_J^* must be sufficiently greater than t^* such that when evaluated at $t = t^*$

$$\sum_{i \in J} [p_i (hw_{iz}^e - q) dz_i^e + (1 - p_i) (hw_{iz} - q) dz_i] > - \sum_{i \in J} [p_i h_{\beta_j} w_i^e + (1 - p_i) h_{\beta_j} w_i].$$

Q.E.D.

References

- Becker, Gary S. and Kevin M. Murphy (1988), "A Theory of Rational Addiction." *Journal of Political Economy*, 96(4), pp. 675-700.
- Becker, Gary S., Kevin M. Murphy and Michael Grossman (2006), "The Market for Illegal Goods: The Case of Drugs." *Journal of Political Economy*, 114(1), pp. 38-60.
- Bernheim, B. Douglas and Antonio Rangel (2004), "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 94(5), pp. 1558-1590.
- Bernheim, B. Douglas and Antonio Rangel (2007), "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision-Makers." Ch. 2 in Peter A. Diamond and Hannu Vartiainen (eds.), *Behavioral Economics and Its Applications*. Princeton University Press.

Choi, James J., David I. Laibson, Brigitte C. Madrian and Andrew Metrick (2003), "Optimal Defaults." *American Economic Review Papers & Proceedings*, 93(2), pp. 180-185.

Gruber, Jonathan and Botond Köszegi (2001), "Is Addiction "Rational"? Theory and Evidence." *Quarterly Journal of Economics*, 116(4), pp. 1261-1303.

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and Self-Control." *Econometrica*, 69, pp. 1403-1436.

Gul, Faruk and Wolfgang Pesendorfer (2004), "Self-control and the Theory of Consumption." *Econometrica*, 72, pp. 119-158.

Gul, Faruk and Wolfgang Pesendorfer (2007), "Harmful Addiction." *Review of Economic Studies*, 74(1), pp. 147-172.

Hansen, Jes Winther (2005), "Present-Biased Individuals, Optimal Paternalism and Transfers in Kind." EPRU Working Paper 05-11.

Laibson, David I. (1997), "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2), pp. 443-477.

Laibson, David I. (2001), "A Cue Theory of Consumption." *Quarterly Journal of Economics*, 116(1), pp. 81-120.

O'Donoghue Ted and Matthew Rabin (1999), "Doing It Now or Later." *American Economic Review*, 89(1), pp. 103-124.

O'Donoghue Ted and Matthew Rabin (2001), "Choice and Procrastination." *Quarterly*

Journal of Economics, 116(1), pp. 121-160.

O'Donoghue Ted and Matthew Rabin (2003), "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review Papers & Proceedings*, 92(2), pp. 186-191.

O'Donoghue Ted and Matthew Rabin (2006), "Optimal Sin Taxes." *Journal of Public Economics*, 90, pp. 1825-1849.

Saez, Emmanuel (2002), "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes." *Journal of Public Economics*, 83, pp. 217-230.

Thaler, Richard H. and Cass R. Sunstein (2003), "Libertarian Paternalism." *American Economic Review Papers & Proceedings*, 93(2), pp. 175-179.

Chapter 5

An Evaluation of the Tax-Transfer Treatment of Married Couples in European Countries*

Herwig Immervoll, OECD and ISER at University of Essex

Henrik Jacobsen Kleven, London School of Economics

Claus Thustrup Kreiner, University of Copenhagen

Nicolaj Verdelin, University of Copenhagen

August 2008

Abstract

This paper presents an evaluation of the tax-transfer treatment of married couples in 15 EU countries using the EUROMOD microsimulation model. First, we show that many tax-transfer schemes in Europe feature negative jointness defined as a situation where the tax rate on one person depends negatively on the earnings of the spouse. This stands in contrast to the previous literature, which has focused on a specific form of positive jointness. The presence of negative jointness is driven by family-based and means-tested transfer programs combined with tax systems that usually feature very little jointness. Second, we consider the labor supply distortion on secondary earners relative to primary earners implied by the current tax-transfer systems, and study the welfare effects of small reforms that change the relative taxation of spouses. We present microsimulations showing that simple revenue-neutral reforms that lower the tax burden on secondary earners are associated with substantial welfare gains in most countries. Finally, we consider the tax-transfer implications of marriage and estimate the so-called marriage penalty. For most countries, we find large marriage penalties at the bottom of the distribution driven primarily by features of the transfer system.

Keywords: Taxation of couples; Negative jointness; Marriage Penalty; Tax reform

JEL classification: H20

*Comments by seminar participants at the CESifo 2008 area conference on public sector economics are gratefully acknowledged. The project has been supported by a grant from the Economic Policy Research Network (EPRN).

1 Introduction

The tax treatment of couples has been a debating point throughout the existence of the income tax. Actual policies have varied over time and across countries. Over the past three decades, there has been an international trend from joint to individual taxation of husbands and wives, and today the majority of OECD countries use the individual as the basic unit of taxation. Under fully individual taxation, tax liability is assessed separately for each family member and is therefore independent of the income of other individuals living in the household. By contrast, in a system of fully joint taxation of couples, as operated by for example the United States, tax liability is assessed at the family level and depends on total family income. Three basic points have been noted in previous discussions of the choice between individual and joint taxation (e.g., Rosen, 1977; Boskin and Sheshinski, 1983; Pechman, 1987).

The first argument is an efficiency argument. It starts from the empirical observation that the secondary earner in a family—typically the wife—tends to have a more elastic labor supply than the primary earner (e.g., Blundell and MaCurdy, 1999). A Ramsey-type optimal tax rule then suggests that the labor income of secondary earners should be taxed at a lower rate than the labor income of primary earners. This is achieved to a certain degree by a progressive individual income tax, because primary earners have higher incomes and therefore face higher marginal tax rates than their spouses. On the other hand, a fully joint income tax creates identical marginal tax rates across members of the same household and hence does not meet this efficiency criterion.

The second argument is that tax systems should be neutral with respect to marriage decisions. This can be viewed as an efficiency argument that tax systems should not distort the marriage market or as a horizontal equity argument that identical couples (married or cohabitating) should be treated identically for tax purposes. While individual-based taxation is neutral with respect to marriage, joint tax systems are generically non-neutral. Jointness may penalize or subsidize marriage depending on the exact design, and the size of penalties/subsidies generally depends on the distribution of income within the family.

The third argument is an equity argument, taking as its point of departure that welfare is better measured by family income than individual income. As a result, two families with the same total income should, other things being the same, pay equal taxes. By the same token, if one family receives a higher total income than another family, then the first family should face a higher tax liability than the second one. This equity criterion is satisfied by a joint income tax that depends on total family income, but not by a progressive individual tax system, because in that case tax liability depends also on the distribution of incomes within the family.

This paper attempts to shed light on the three issues discussed above. We start by noting that these issues ultimately pertain to the redistribution scheme as a whole, not just the tax system, and we therefore present an integrated treatment of the tax and transfer system. A recurrent theme in the paper is that the transfer system is a crucial element in understanding and evaluating redistribution schemes affecting married couples. We also

point out that the focus in previous discussions on the choice between individual tax treatment and joint tax treatment based on family income represents an oversimplification, because real-world redistribution schemes are almost never fully individual or fully joint. There are two reasons for this. First, while most countries have adopted individual filing in the tax system, they tend to retain certain elements of jointness such as the transfer of unused allowances across spouses, dependent spouse exemptions, etc. Second, transfer systems are nearly always fully joint, because social benefits are means-tested according to the combined income and assets of the two spouses in the household. This implies that actual redistribution systems typically combine a form of quasi-individual taxation with family-based transfer systems, creating a fairly complicated jointness structure that is different from the two polar cases typically analyzed.

The paper presents a comprehensive evaluation of the tax-transfer treatment of married couples in 15 EU countries. The analysis has three components. First, we carefully map the precise nature of jointness in tax-transfer schemes in our sample of countries using the EUROMOD microsimulation model. EUROMOD is built around country-specific, but partly harmonized, micro datasets combined with a detailed tax-benefit simulator capturing the full set of institutional features of tax and transfer systems in each country.¹ We find that many tax-transfer schemes in Europe feature *negative jointness* defined as a situation where the tax rate on one person depends negatively on the earnings of the spouse. Such a system is opposite to the form of jointness typically

¹An introduction to EUROMOD and a descriptive analysis of taxes and transfers in the EU countries has been provided by Sutherland (2001), Immervoll and O'Donoghue (2003), and Immervoll (2004).

analyzed in the academic literature—fully joint and progressive taxation—because such schemes feature a positive interaction between tax rates and spousal earnings and therefore *positive jointness*. The presence of negative jointness is driven by the interaction between family-based transfers and individual or almost-individual taxes. To see this, consider a secondary earner, say the wife, deciding about labor market entry. If she is married to a low-income husband, the family is in the phase-out range of transfer programs, and she will face a high effective tax rate. On the other hand, if she is married to a high-income husband, the family is beyond the phase-out range of transfer programs, and she will face a low effective tax rate because the income tax is individual. Hence, the wife’s tax is declining in the husband’s earnings.

Second, the paper considers the incentives to supply labor for secondary earners relative to primary earners implied by the existing tax-transfer systems, and studies the welfare effects of reforms that change the relative taxation of spouses. This issue is separate from the nature of jointness discussed above: jointness has to do with the relationship between tax rates and *spousal* earnings (a cross-derivative in the tax function), whereas labor supply incentives have to do with the relationship between tax liability and *own* earnings (an own-derivative in the tax function). Previous work has often discussed the two issues as if they are one and the same, but we find that the distinction is empirically important. To study the welfare effects of tax-transfer reform, the paper starts by setting out a simple theoretical model that incorporates only participation responses, and then turn to a general model that allows for both participation and hours-of-work responses

for both spouses in the household. Microsimulations of different revenue-neutral reforms that reduce the tax burden on secondary earners show that, for both models and for most countries, a lowering of the tax burden on secondary earners is associated with substantial welfare gains.

This part of the paper may be seen as an extension of our previous work based on single-person households (Immervoll et al., 2007) to the case of two-person households. It is also related to the recent work by Alesina and Ichino (2007), arguing that tax schemes should be gender-specific with lower rates on females. We do not consider gender-specific taxation as such (consistent with real-world tax systems that are anonymous and hence gender-blind), but consider reforms that change the taxation of primary versus secondary earners. We define primary versus secondary, not in terms of gender, but in terms of relative earnings within the family—a concept that is correlated with gender.² Indeed, it is shown that in almost all countries more than 80% of secondary earners are women, and in some countries more than 90% of secondary earners are women. Thus, the reforms under consideration strongly target married women with low earnings or weak labor market attachment without formally discriminating based on gender. This is important because gender-specific taxation per se would be unconstitutional in most countries.

Third, the paper explores the distortions in the decision to marry by simulating the size of marriage penalties resulting from the combined effect of taxes and transfers. The presence of fully family-based transfers generally penalizes marriages at the bottom of

²More specifically, the lower-earning spouse in each family is defined as the secondary earner. For one-earner couples, this obviously implies that the non-working spouse is the secondary earner.

the distribution, implying that marriage penalties at the bottom in practice go hand in hand with negative jointness.³ Indeed, we find large marriage penalties at the bottom of the distribution (but not at the top) in most countries, which raises important questions pertaining to fairness as well as to efficiency. Transfers and taxes that depend on marriage are often accused by conservatives of destroying the traditional two-parent family and leading to high rates of single motherhood. Although empirical studies of the effects on marriage and divorce from income taxes (e.g. Alm and Whittington, 1997, 1999), welfare benefits (e.g., Hoynes 1997a,b; Moffitt, 1998), or taxes and benefits combined (Dickert-Conlin, 1999; Eissa and Hoynes, 2000b) tend to find modest or no effects, the existence of marriage disincentives continues to be a controversial point of contention.

Most of the literature studying the optimal design of tax and transfer programs and the evaluation of tax and welfare reform rests on models of single-person households. However, real-world tax-transfer schemes for a large part redistribute income across families formed around couples, creating a substantial gap between theory and practice. This has triggered a recent and growing interest in generalizing the theory of optimal income redistribution to explicitly deal with couples. For example, Kleven, Kreiner, and Saez (2007, 2008) explore the optimal nonlinear taxation of couples as a multi-dimensional screening problem, whereby agents (couples) are characterized by a multi-dimensional parameter (ability and taste-for-work parameters of each spouse) that are unobserved by the principal (the government which maximizes social welfare). They find that, under certain

³However, theoretically it is entirely possible to design a negatively joint tax system that subsidizes rather than penalizes marriage.

assumptions, optimal incentive schemes feature negative jointness, which is consistent with our findings for Europe. Recent papers by Brett (2006) and Cremer, Lozachmeur and Pestieau (2006) also analyze the optimal taxation of couples as a multidimensional screening problem. The rest of the literature (e.g. Schroyen, 2003; Alesina and Ichino, 2007; Apps and Rees, 2007) typically restricts the tax function to be separable (albeit gender-specific), thereby sidestepping the complexities associated with multi-dimensional screening.

Our paper may be seen as an applied counterpart to these recent theoretical papers. By focusing on small reforms rather than the optimal system, we are able to set out a tractable analysis based on more general and realistic labor supply models than the very stylized models previously considered.

The paper is organized as follows. Section 2 briefly describes the data and the EURO-MOD model. Section 3 maps out the existing tax-transfer treatment of married couples in our sample of European countries. Section 4 sets out a joint labor supply model to evaluate reforms that reduce the tax burden on second-earner participation, and goes on to present a microsimulation study of specific reforms. Section 5 studies marriage penalties, while Section 6 offers concluding remarks.

2 Data

Our data source is EUROMOD, a microsimulation model for the EU built around partly homogenized micro datasets that include data on earnings, labor force participation and

demographics. The version available for this study relates to 1998 and covers the 15 countries that constituted the EU at that time. Based on detailed algorithms capturing the full range of institutional features of tax and transfer systems in each country, the model is able to compute a wide range of taxes and benefits for each observation unit in representative samples for the various countries. The main policy instruments incorporated in EUROMOD are income taxes, social security contributions (or payroll taxes) paid by employees, benefit recipients, and employers as well as universal and means-tested social benefits including housing assistance.⁴ The model fully accounts for the complicated interaction of different types of taxes and benefits with earnings, assets, employment status, marital status, housing situation and children, and its considerable level of detail makes it an ideal tool for comparative tax analysis.⁵

We restrict the sample to married couples where both husband and wife are between 16 and 64 years of age, where the couple as a whole reports positive annual earnings, and where at least one member of the household has been working the entire year. We exclude those who are currently receiving pension, early retirement, or disability benefits. In each couple, we define the primary earner (PE) as the highest-earning member and the secondary earner (SE) as the lowest-earnings member of the household. Together

⁴In the results reported here, we do not include unemployment insurance (UI) benefits in the calculation of effective participation tax rates. This is due in part to difficulties associated with accounting properly for the implications of limited UI duration in our static tax rate measures. At a more conceptual level, it is likely that UI schemes providing insurance against involuntary and temporary job loss have very different incentive implications than poverty alleviation programs offering permanent income guarantees to all non-workers.

⁵For further information on EUROMOD, the reader is referred to Sutherland (2001) as well as the Internet at <http://www.iser.essex.ac.uk/msu/emod>

with our sample restriction, this implies that, in one-earner couples, the primary earner works the entire year while the secondary earner is non-employed throughout the year. In two-earner couples, the secondary earner works either part of or the entire year but always has relatively low earnings.

While we feel that it makes sense to define primary versus secondary earner in terms of earnings (and indirectly labor market participation), the earnings-based definition is in practice highly correlated with a gender-based definition. To demonstrate this, Table 1 displays the share of women among secondary earners according to our definition. We see that, in one-earner couples, more than 90 percent of secondary earners (non-participants) are women in all countries except the Nordic countries and the United Kingdom. In two-earner couples, the second-earner definition is slightly less skewed towards women, so that on average the share of secondary earners that are women varies between 80 and 90 percent across most of the 15 countries in our sample. The close relationship between relative earnings within families and gender implies that a purely earnings-based couple tax function can be targeted to gender without being formally gender-based. This is important because gender-specific taxation per se would be unconstitutional in most countries.⁶

⁶Despite the economic equivalence between gender-based taxation and affirmative action (which has been approved by the courts on many occasions), the two policies would be viewed very differently by the courts. See Rubinfeld (1997) for a discussion of this point in the context of race.

Table 1. Share of women among secondary earners

Country	One-earner couples	Two-earner couples	One- and two-earner couples
Austria	0.96	0.87	0.90
Belgium	0.92	0.73	0.78
Denmark	0.83	0.82	0.82
Finland	0.65	0.76	0.75
France	0.93	0.74	0.80
Germany	0.91	0.78	0.83
Greece	0.98	0.75	0.90
Ireland	0.96	0.78	0.89
Italy	0.94	0.81	0.88
Luxembourg	0.99	0.77	0.89
Netherlands	0.95	0.88	0.91
Portugal	0.95	0.78	0.85
Spain	0.97	0.73	0.89
Sweden	0.66	0.76	0.75
United Kingdom	0.78	0.83	0.81

Note: Secondary earners are defined as the spouses with the lowest earnings in the couples. Source: EUROMOD Microsimulation Model.

3 The Tax-Transfer Treatment of Couples in Europe

Based on EUROMOD, this section maps out the tax-transfer treatment of married couples in our sample of European countries. Table A3 in the appendix summarizes the most important institutional features of tax and benefit systems affecting married couples in each country.

It is useful to start by distinguishing between two different properties of a tax-transfer scheme for couples: (i) the relative tax rates on husbands and wives, and (ii) the jointness of the tax schedule. In most of the existing literature, these two properties have been studied as if they are one and the same, with joint taxation being defined as a situation with identical marginal tax rates on the two spouses and individual taxation being defined as a situation with a higher marginal tax rate on the primary earner. However, this close

relationship between relative tax rates and jointness is present only under very strong restrictions on the tax schedule. In general, jointness is related to the cross-relationship between tax rates and *spousal* earnings (a cross-derivative in the tax function), whereas tax rates reflect the relationship between tax liability and *own* earnings (an own-derivative in the tax function). In principle and in practice, it is entirely possible to combine forms of jointness with, say, lower tax rates on secondary earners.

To make the discussion precise, it is helpful to define the tax function for couples as $T(z_p, z_s)$, where z_p denotes primary earnings and z_s denotes secondary earnings. Below we often refer to this as a ‘tax function’, but we want to think of $T(\cdot)$ as the net payment by a couple to the government embodying taxes as well as transfers. The effective marginal tax rate (including benefit phase-out) on each spouse is given by $T'_p(z_p, z_s)$ and $T'_s(z_p, z_s)$. Marginal tax rates are of course important for determining hours worked for those who are working (the intensive margin of labor supply). We report marginal tax rates in Appendix A for the interested reader, but focus instead on a different tax rate measure—the participation tax rate. This is a more interesting tax rate measure because the extensive margin of labor supply is empirically more important. We define the participation tax rate on a particular family member as the total change in $T(\cdot)$ as this family member enters into employment as a share of earnings generated by the entry. In order to calculate participation tax rates, one has to make assumptions about the sequence of participation choices within the household because, with jointness in the tax-transfer code, the tax liability change associated with a person entering depends on

whether the spouse is working or not. We make the natural assumption that the primary earner enters first and the secondary earner enters second. In Section 4.1, we provide a microfoundation for this model, which has been adopted in many empirical labor supply studies (e.g., Eissa, 1995; Eissa and Hoynes, 2004).

Under the assumed sequence of labor market entries in the household, the participation tax rates on the primary and secondary earners are given by

$$\tau_p \equiv \frac{T(z_p, 0) - T(0, 0)}{z_p}, \quad \tau_s \equiv \frac{T(z_p, z_s) - T(z_p, 0)}{z_s}. \quad (1)$$

These tax rates are simulated by EUROMOD in the following way. For the computation of τ_s , we consider the subsample of two-earner couples and start by computing actual taxes net of transfers $T(z_p, z_s)$ at each observed earnings pair, accounting for other relevant household information (place of residence, number of kids, etc.). We then recompute taxes and transfers in the alternative (hypothetical) situation where the secondary earner does not work, $T(z_p, 0)$, and calculate τ_s as in eq. (1). Analogously for τ_p , we use the sample of one-earner couples to simulate taxes net of transfers in the original situation, $T(z_p, 0)$, and in the alternative situation where the primary earner is not working, $T(0, 0)$, and then apply formula (1).⁷

Table 2 shows participation tax rates and labor market outcomes for primary and

⁷Our tax-rate estimates are therefore calculated for those currently working. As a result of sample selection, one would expect tax rates to be different for non-working individuals considering a transition into work. As we do not observe the earnings potential of non-working individuals, calculating their participation tax rates would require jointly estimating a wage and participation model for couples. In the microsimulation exercise in Section 4, we deal with the selection issue indirectly by considering a decreasing profile of participation elasticities such that new labor market entrants tend to be located at the bottom end of the income distribution.

secondary earners in each country (averages for each country sample). As one would expect, Scandinavia and Northern-Continental Europe feature higher overall tax rate levels than Anglo-Saxon and Southern European countries. More interestingly, the tax rate on primary earners is higher than on secondary earners in all but the four Southern European countries (Greece, Italy, Portugal, and Spain). This is a result of the impact of family-based and means-tested welfare benefits, which are affected more by the first than by the second entrant. We do not observe the same effect in Southern Europe where welfare benefits are less generous. Although most countries impose a higher participation tax rate on the primary earner, there are substantial differences in the relative rates across countries. In particular, the UK system stands out by being much more favorable to second-earner participation than all other countries.

Table 2. Participation tax rates and labor market outcomes

Country	Primary earners participation tax	Secondary earners participation tax	PE tax / SE tax	Relative participation (PE / SE)	Relative earnings (PE / SE)
Austria	0.63	0.36	1.72	1.66	2.18
Belgium	0.73	0.74	0.98	1.38	1.95
Denmark	0.73	0.53	1.38	1.13	1.74
Finland	0.60	0.36	1.65	1.05	1.66
France	0.85	0.63	1.35	1.53	1.67
Germany	0.63	0.51	1.22	1.53	1.99
Greece	0.27	0.28	0.97	3.13	1.61
Ireland	0.54	0.44	1.22	2.55	2.18
Italy	0.35	0.46	0.77	2.15	1.51
Luxembourg	0.50	0.32	1.54	2.52	2.25
Netherlands	0.56	0.44	1.28	1.68	2.61
Portugal	0.37	0.41	0.90	1.74	1.56
Spain	0.34	0.41	0.83	3.25	1.61
Sweden	0.66	0.51	1.28	1.07	1.53
United Kingdom	0.56	0.20	2.79	1.47	2.20

Note: The first two columns list the average effective participation tax rates for primary earners in one-earner couples and secondary earners in two-earner families, respectively. Column 4 shows relative participation rates of primary and secondary earners and column 5 lists the relative average earnings of primary and secondary earners conditional on working. The calculation of the tax rates is described in the text. Source: EUROMOD Microsimulation Model.

The ratio of the primary-earner tax rate to the secondary-earner tax rate is interesting

because it can be compared to optimal tax rules expressing relative tax rates as a function of elasticities (e.g., Boskin and Sheshinski, 1983; Kleven and Kreiner, 2006). In the special case of separability in utility of spousal labor supplies, the optimal tax rate on each spouse is given by an inverse elasticity rule and the optimal relative tax rate τ_p/τ_s is therefore given by the participation elasticity of the secondary earner relative to the primary earner. This implies that the tax ratios in the table can be seen as critical values for relative participation elasticities. For example, in the United Kingdom, if the second-earner elasticity is more than 2.79 times as high as the primary earner elasticity, it would be efficient to shift some of the tax burden from secondary earners to primary earners. In view of the evidence on the responsiveness of labor force participation of married women, the table seems to suggest that in many countries the relative tax rate on secondary earners is inefficiently high. We return to this issue in Section 4.

Finally, the table shows that both participation and earnings (conditional on participation) tend to be strongly skewed in favor of primary earners in most countries. Although the countries we consider differ along many dimensions (besides tax rates) that may have direct implications for labor market outcomes, it is interesting that the cross-country variation in relative participation rates is somewhat consistent with the variation in relative participation taxes. For example, Southern European countries are characterized by lower participation taxes along with higher participation rates for primary relative to secondary earners compared to most other countries. At the other end of the spectrum, Denmark, Finland, and the UK are associated with higher relative tax rates

and lower relative participation for primary earners.

Let us now consider the jointness of the couple tax function $T(z_p, z_s)$ and therefore the cross-derivative T''_{ps} . One benchmark case is that of fully individual taxation, i.e. $T(z_p, z_s) = T_p(z_p) + T_s(z_s)$, which is associated with a zero cross-derivative $T''_{ps} = 0$. In practice, the functional forms $T_p(\cdot)$ and $T_s(\cdot)$ would typically be the same, in which case we have a so-called *anonymous* individual tax. Another benchmark case is the fully joint couple tax $T(z_p + z_s)$ as adopted in the United States and in some European countries. If the system additionally features a progressive marginal tax rate structure ($T'' > 0$), the couple tax would be associated with a positive cross-derivative $T''_{ps} > 0$. More generally, there is a whole range of joint couple tax functions $T(z_p, z_s)$ with $T''_{ps} \neq 0$. Kleven, Kreiner, and Saez (2007) define *positive jointness* as a system where the tax on one person depends positively on spousal earnings ($T''_{ps} > 0$), and *negative jointness* as a system where the tax on each partner depends negatively on spousal earnings ($T''_{ps} < 0$). Because a fully joint schedule is associated with positive jointness, individual taxation can be seen as an intermediate (rather than polar) case in between full jointness and negative jointness. This is interesting because we show below that many real-world schedules feature negative jointness, implying that they have moved further away from fully joint taxation than the individual system.

The above definitions of jointness are stated in terms of cross-derivatives of marginal tax rates. Consistent with the analysis of tax rate levels, we will state a definition of jointness in terms of participation tax rates. In particular, we say that a system is

positively joint if $\frac{\partial \tau_s}{\partial z_p} > 0$, negatively joint if $\frac{\partial \tau_s}{\partial z_p} < 0$, and separate if $\frac{\partial \tau_s}{\partial z_p} = 0$. While the definitions of jointness in terms of marginal tax rates are local, the definitions in terms of participation (‘average’) tax rates reflect that a tax schedule is joint on average over a range of incomes.⁸

Before turning to the empirical results, it is helpful for a moment to separate the tax and transfer system. Denote by $t(z_p, z_s)$ the tax payment and by $b(z_p, z_s)$ the benefit payment, so that $T(z_p, z_s) = t(z_p, z_s) - b(z_p, z_s)$. Consider then a tax-transfer scheme that combines an individual income tax and a fully joint transfer system, i.e. $T(z_p, z_s) = t_p(z_p) + t_s(z_s) - b(z_p + z_s)$. The definition in eq. (1) then implies $\tau_s = [t_s(z_s) + b(z_p) - b(z_p + z_s)] / z_s$. Means-testing corresponds to $b(z_p) - b(z_p + z_s) \geq 0$, which creates an extra tax on second-earner participation. However, as z_p increases, the family is pushed beyond the phase-out range of the various transfer programs (at any given z_s), which tends to lower $b(z_p) - b(z_p + z_s)$ and create a pattern where $\frac{\partial \tau_s}{\partial z_p} < 0$ at the bottom. This explains a pattern we find for many countries.

For our measurement of jointness, we construct a number of hypothetical households that vary with respect to household earnings and the number of children, and apply EUROMOD to calculate effective tax rates for these hypothetical families.⁹ We base this part of the analysis on hypothetical households (instead of data on actual households)

⁸In Appendix A, we present empirical measures of jointness based on marginal tax rates. The qualitative results are similar to those presented below, but participation tax rate measures are smoother because they reflect average jointness of marginal tax rates over a range of incomes.

⁹Because we are working with hypothetical households, it is necessary to make an assumption about the living arrangements of the families. We have chosen to assume that all families reside in rental housing, and have then imputed rental costs for all countries.

in order to adequately isolate the interdependence between spouses in the tax-benefit legislation. If we were to use sample data and compare the net-tax burden of actual households at different earnings levels, the results would be affected by selection effects.¹⁰

To illustrate the jointness in the tax-transfer system, we plot the participation tax rate of married individuals at different income levels as a function of the earnings of the spouse. We consider married individuals at four different income levels: the 5th, 10th, 50th, and 90th percentiles of the earnings distribution of secondary earners (denoted below by SEp5, SEp10, SEp50, and SEp90). For each of these individual income levels, we calculate the participation tax rate at 20 different earnings levels of the spouse, corresponding to the 5th, 10th, 15th,..., 100th percentiles in the earnings distribution of primary earners (denoted below by PEp5, PEp10, ..., PEp100). Our results are shown in Figure 1 for families with two children. A corresponding figure for childless couples is provided in Appendix B.¹¹

The most striking result is that most countries display substantial negative jointness at the bottom of the income distribution. As explained above, this can be largely attributed to means-tested benefits such as social assistance, housing benefits and child benefits that are phased-out as a function of total household income.¹² Indeed, the high claw-back

¹⁰For instance, marriage patterns are known to display positive assortative matching, which in itself would tend to produce a positive relationship between individual tax rates and spousal earnings.

¹¹When considering a couple with one spouse belonging to the bottom of the earnings distribution for primary earners and the other spouse located at the top of the earnings distribution for secondary earners, it may actually be the second spouse who has the highest earnings. This is, however, only relevant for the lower part of the two grey curves because the earnings of the secondary earners in the data are substantially lower than primary earnings. Moreover, the slopes of the curves still reveal the type of jointness at these earnings combinations.

¹²In Germany and Belgium, there is an initial slight increase in the second-earner tax rate at low levels of primary-earner income provided that the secondary earner also enters at a low earnings level. This is

Figure 1. Participation tax rates of secondary earners (in percent) at selected earnings levels as a function of primary earnings (vingtiles) for couples with two children

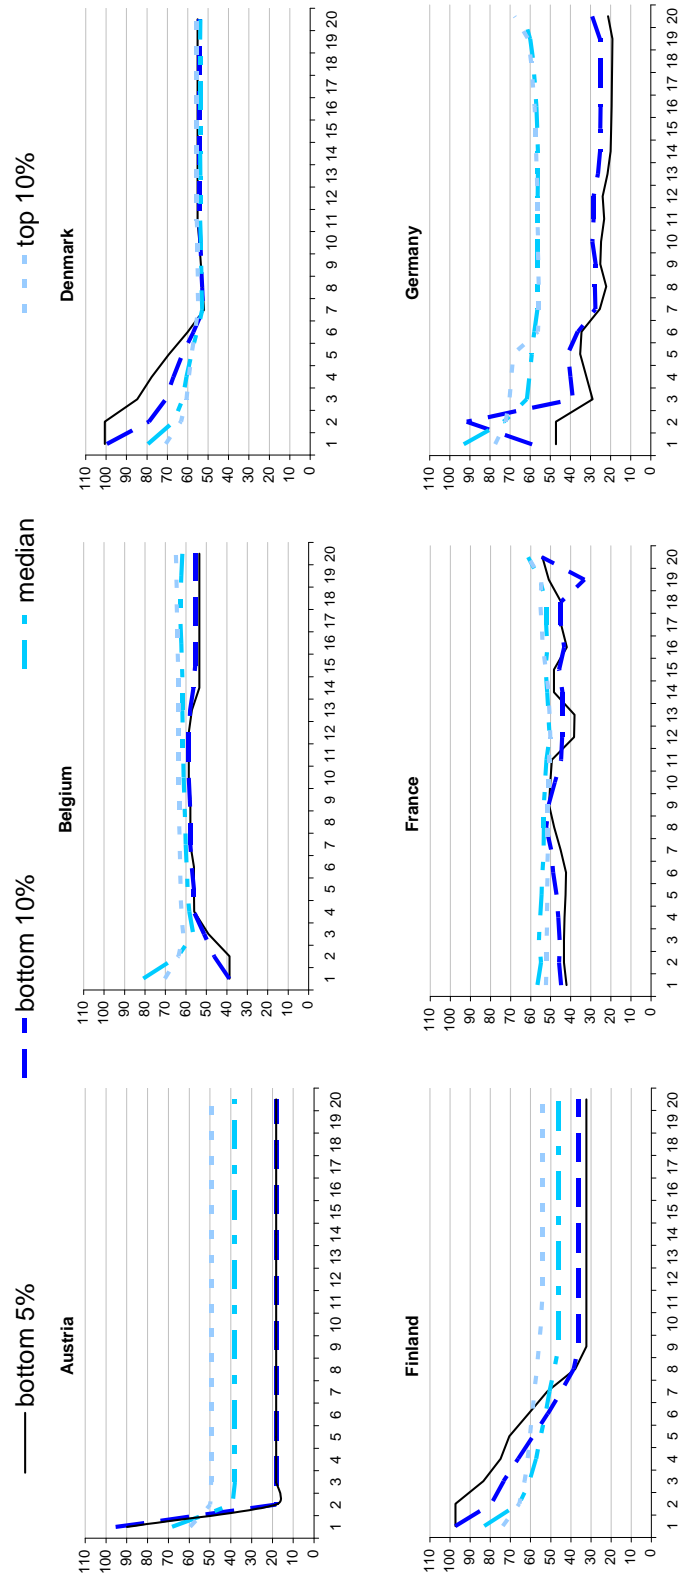
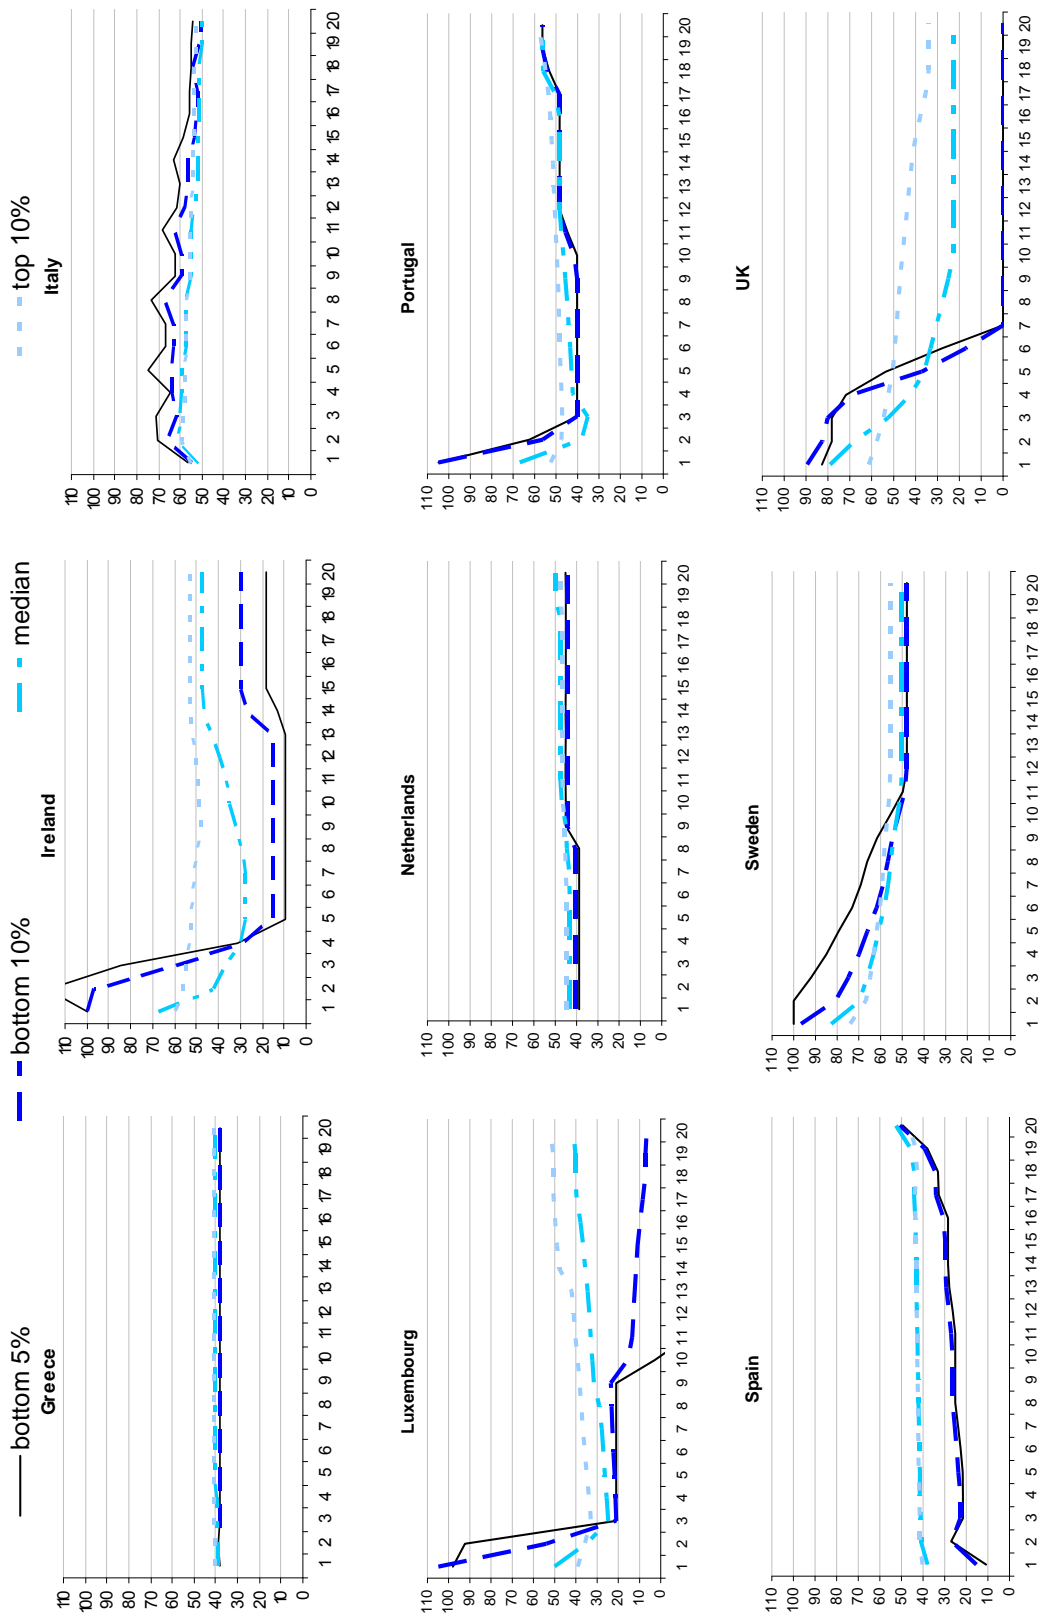


Figure continues on the next page

Figure 1 cont.



Note: Each figure shows the participation tax rate of the secondary earner (SE) in percent as a function of the earnings of the primary earner (PE), depicted by vingtile group in the PE earnings distribution. The graphs are shown for four different SE earnings levels: the lowest vingtile (SEp5), the lowest decile (SEp10), the median (SEp50), and the top decile (SEp90) of the SE earnings distribution. Source: EUROMOD microsimulation model.

rates used in many countries tend to generate participation taxes that are very high for secondary earners married to low-wage primary earners, often above 70% and sometimes close to 100%.

Countries with negative jointness at low income levels may be divided into two groups depending on the pattern at higher income levels. Countries that operate an individual income tax (possibly apart from some family-based tax expenditures at the bottom) and/or have a fairly flat tax rate structure at the top tend to converge towards no jointness as the income of the primary earner becomes high.¹³ Austria, Denmark, Finland, Sweden and the United Kingdom display this pattern, which we may label *negative-neutral* jointness. The strongest example of this pattern is perhaps the United Kingdom where negative jointness at the bottom is driven by both the welfare and the tax system. The Working Families Tax Credit (WFTC), an in-work cash benefit provided through the tax system, is based on household income and is phased out at a rate of 70%. The combination of the WFTC and the withdrawal of housing benefits and social assistance creates participation taxes on secondary earners married to low-earning spouses around 70-90%. While second-earner labor market entry in the UK is strongly discouraged in low-income families, it is encouraged in higher-income families due to the individual income tax system. In particular, because working spouses with low earnings are not

because these countries employ an earnings disregard in the transfer system, so that the lowest-income families are not affected by benefit withdrawal.

¹³Notice that a flat (linear) income tax, even if it is based on family income, effectively implies separability in the tax treatment of spouses. As an example, this is important for Denmark, which operates a form of joint taxation by combining individual filing with the possibility of transferring certain allowances and exemptions across spouses. However, because the marginal tax rate structure is quite flat there is very little jointness at the top.

liable to pay either income tax and social insurance contributions, the second-earner tax rate at SEp5 and SEp10 drops to zero once primary-earner income exceeds PEp30 and stays at zero as primary earnings increase.

Another group of countries combine negative jointness at the bottom with positive jointness at the top. Countries with this pattern of *negative-positive* jointness are Belgium, France, Germany, Ireland, Luxembourg and Portugal. All of these countries operate a progressive tax system based on family income causing the secondary-earner participation tax to be increasing in primary-earner income once the family is beyond the phase-out range of welfare programs. However, the degree of positive jointness at the top is generally quite weak and much less salient than the negative jointness at the bottom. This may seem surprising but can be explained by the fact that the marginal tax rate structure is quite flat in most European countries, partly as a result of upper contribution limits built into social security contribution schedules. Indeed, as discussed above, a completely linear tax system, even if it is based on family income, effectively implies separability in the tax treatment of spouses. Notice also that, in France, the curve is relatively flat both at the bottom and at the top because the withdrawal of various family benefits and housing assistance occurs at different income levels and tends to offset the presence of positive jointness in the income tax.¹⁴

Greece, Italy, and to some extent the Netherlands are the only countries that show

¹⁴In France, the drop in the participation tax rate for low-wage secondary earners (at SEp10) when the primary-earner income becomes very high (at PEp95) is due to complicated features of the income test for family benefits (*Allocations Familiales*) that were in place only in 1998, the year of our sample.

virtually *no jointness*. All three countries operate individual income tax systems, and in Greece and Italy only very limited means-tested benefits are available.¹⁵ The Netherlands does offer family-based social assistance, but primary earnings are higher in the Netherlands than in most other countries, implying that transfer phase-out plays a limited role for second-earner labor market entry.¹⁶ Spain is the only country characterized by *positive jointness*. There is no social assistance and the design of the Spanish income tax implies that, for low-income families, it is optimal to file under the optional joint tax. For higher-income families, it is typically optimal to file separately, which explains why there is less jointness if the secondary earner is at the median or above.

4 A Welfare Evaluation of Cutting Taxes for Secondary Earners

It is often argued that the tax burden on secondary earners should be reduced in order to increase economic efficiency. Indeed, a traditional Ramsey-type efficiency argument calls for a low marginal tax rate on secondary earners because their labor supply is relatively elastic (Rosen, 1977; Boskin and Sheshinski, 1983). The traditional argument is derived in a model with only hours-of-work responses and where the tax system is restricted to be linear (albeit gender specific). However, the modern empirical labor supply literature

¹⁵In Greece, no means-tested benefits are available for married couples. In Italy, such benefits are very limited, especially for couples without kids as reflected by the almost completely flat curve (in appendix) for those couples. Further, family benefits in Italy are phased-out in discrete amounts at different income levels, which accounts for the small bumps visible for low-income families.

¹⁶The small bump around PEP40 reflects mandatory health insurance contributions for non-working spouses that apply to primary earners with earnings below a certain threshold. Above this earnings level, health insurance contributions for non-working spouses are voluntary and hence not counted as taxes.

shows that the strong responsiveness of the labor supply of married females is driven by labor force participation, not by hours worked for those who are working (e.g., Heckman, 1993; Blundell and MaCurdy, 1999). This calls for a policy that reduces the participation tax rate on secondary earners.

A policy change should be evaluated, not just in terms of efficiency, but also with respect to its consequences for distributional equity. A revenue-neutral reform reducing the participation tax rate on secondary earners necessarily implies a redistribution in favor of two-earner couples at the expense of one-earner and/or zero-earner couples. To the extent that two-earner couples are better off than zero- and one-earner couples such reforms come at the cost of a reduction in distributional equity. While the statement that two-earner couples are better off than zero-earner couples seems noncontroversial, the comparison between one- and two-earner couples is less clear-cut. Notice first that, for a given level of primary earnings, the notion that two-earner couples are better off than one-earner couples is consistent with the underlying assumption in all of the optimal income tax literature that higher household income is a signal of higher utility.¹⁷ Whether two-earner couples are better off *on average* depends also on the sorting in the marriage market. Positive sorting in earnings (such that two-earner couples tend to have higher primary-earner income along with the presence of secondary-earner income) reinforces the view that two-earner couples are better off. On the other hand, if there is

¹⁷In the presence of general non-linear tax instruments, the relevant comparison for the determination of the optimal tax on secondary entry is indeed between different types of couples at a given level of primary earnings (Kleven, Kreiner, and Saez, 2007).

negative sorting whereby rich people tend to have non-working spouses, it is theoretically possible that one-earner couples are better off on average. However, as shown by Kleven, Kreiner, and Saez (2008) for the UK, there is a positive correlation in spousal earnings (conditional on working) combined with a very weak correlation between primary-earner income and spousal labor force participation. All this suggests that two-earner couples are better off, so that lowering the participation tax on secondary earners comes at a cost of distributional equity.

In this section, we start by setting out a simple extensive labor supply model allowing us to evaluate the efficiency-equity trade-off for reforms aimed at increasing second-earner participation. In particular, we consider small (marginal) tax reforms, which provide a transfer to two-earner couples financed by either a tax on both zero- and one-earner couples or a tax on one-earner couples only. The taxes and transfers implemented by the reforms are lump sum conditional on family participation status and therefore do not affect marginal tax rates. Reforms of this type could be implemented in practice by, for example, changing the structure of family allowances. At the end of the section, we generalize the labor supply model to incorporate both intensive and extensive responses for both spouses, and consider reforms that reduce the tax burden on secondary earners by changing marginal tax rates.

4.1 A Simple Joint Labor Supply Model

We consider couples where each spouse decides whether or not to work, but where hours worked (and earnings) conditional on working are fixed. Labor force participation varies

across couples due to heterogeneity in earnings potential and work costs, and households can be grouped into three different categories: no-earner, one-earner, and two-earner couples. In each household, we identify a primary earner and a secondary earner where, by definition, the primary earner enters the labor market first and has higher earnings conditional on working. Each spouse is characterized by a fixed earnings potential, which we denote by (z_p^h, z_s^h) for the two spouses in a household of type h . Letting z_i ($i = p, s$) denote the actual earnings choice, the participation choice for each spouse then amounts to choosing between $z_i = 0$ and $z_i = z_i^h$. The number of households of type h is denoted N_h , $h = 1, \dots, H$, and the total population of households equals $N \equiv \sum_{h=1}^H N_h$.

All households share a common quasi-linear utility function given by

$$u(c, z_p, z_s) = c - q_p \cdot 1(z_p > 0) - q_s \cdot 1(z_s > 0), \quad (2)$$

where c is household consumption, and q_p, q_s denote work costs for the primary and secondary earner, respectively. The work costs capture all costs associated with labor market entry such as a distaste for participation, the value of lost home production, costs of child care and commuting, etc. The indicator function $1(\cdot)$ takes on the value 1 when a given spouse works ($z_i > 0$, $i = p, s$) and zero otherwise. The above utility specification rules out income effects which simplifies considerably the theoretical analysis (Kleven, Kreiner, and Saez, 2007, 2008) as well as the welfare aggregation.

The household faces a non-linear income tax schedule $T(z_p, z_s, \theta)$, where θ is a shift parameter that we use below to capture the effects of a tax reform. The tax function constitutes a net payment to the public sector, embodying both taxes and transfers. The

consumption of each household equals their total net-of-tax earnings, such that eq. (2) can be written as

$$u = z_p + z_s - T(z_p, z_s, \theta) - q_p \cdot 1(z_p > 0) - q_s \cdot 1(z_s > 0). \quad (3)$$

Households choose earnings z_p and z_s so as to maximize eq. (3). For households of type h (i.e., earnings pair z_p^h, z_s^h), there is a distribution of fixed costs described by a continuous joint density function $f_h(q_p, q_s)$ defined over $[0, \infty) \times [0, \infty)$. We define the unconditional density and distribution functions of q_p as $f_h(q_p)$ and $F_h(q_p)$, and the conditional density and distribution functions of q_s as $p_h(q_s | q_p)$ and $P(q_s | q_p)$, and hence the joint density can be written as $f_h(q_p, q_s) = p_h(q_s | q_p) \cdot f_h(q_p)$.

Consistent with much empirical work in this area, we consider households making a sequential labor force participation decision. First, it is decided whether or not the primary earner should enter the labor market and then, conditional on primary-earner participation, it is decided if the secondary earner should join the labor force as well. We need to ensure that the assumed entry sequence is consistent with household optimization, which amounts to a restriction on the joint distribution of fixed work costs. Figure 2 illustrates the problem by depicting the possible joint labor supply choices of the two spouses. The crucial assumption we make is that, both before and after a tax reform, couples are observed only in the shaded areas (0, 1 and 2). The part of the assumption that concerns the initial (before-reform) distribution of couples is innocuous, because we simply define the primary earner (i.e., the highest-earning spouse) in such a way that it is consistent with the permissible pattern. However, when we consider tax reforms




Figure 2: The double-deviation problem.

that induce families to change participation status, we must make sure that no families move to region \emptyset in the figure. This problem is reminiscent of the double-deviation problem in optimal multi-dimensional pricing theory (e.g. Armstrong and Rochet, 1999) and in the theory of optimal taxation with more than one dimension of unobserved household characteristics (e.g., Mirrlees, 1976, 1986; Kleven et al., 2007). While the double-deviation problem poses considerable complexity for studies that attempt to solve for the *optimal* incentive scheme in a multi-dimensional screening context, it is easier to deal with the issue here because we consider only small perturbations (marginal reforms) around an initial equilibrium. Appendix C shows how we deal with the double-deviation issue by imposing restrictions on the distribution of fixed costs.

Given the assumed sequence of labor market entries, a primary earner decides to enter if the net household utility gain of doing so, conditional on spousal non-participation, is

positive. For household h , this implies

$$q_p \leq z_p^h - [T(z_p^h, 0, \theta) - T(0, 0, \theta)] \equiv \bar{q}_p^h, \quad (4)$$

where \bar{q}_p^h is the net-of-tax income gain of primary-earner entry for household type h .

Primary earners with $q_p \leq \bar{q}_p^h$ decide to enter the labor market at $z_p = z_p^h$, whereas primary earners with $q_p > \bar{q}_p^h$ stay outside the labor force. Conditional on primary-earner entry, the secondary earner in household h enters if

$$q_s \leq z_s^h - [T(z_p^h, z_s^h, \theta) - T(z_p^h, 0, \theta)] \equiv \bar{q}_s^h, \quad (5)$$

where \bar{q}_s^h is the net-of-tax income gain of second-earner entry.

Let $E_0^h = N_h [1 - F_h(\bar{q}_p^h)]$, $E_1^h = N_h F_h(\bar{q}_p^h) - E_2^h$ and $E_2^h = N_h \int_0^{\bar{q}_p^h} P_h(\bar{q}_s^h | q_p) f_h(q_p) dq_p$ denote, respectively, the number of zero-earner, one-earner, and two-earner couples of type h . Consistent with our assumed sequence of labor market entry, we define the participation elasticities for primary and secondary earners as

$$\eta_p^h \equiv \frac{\partial E_1^h}{\partial [z_p^h (1 - a_p^h)]} \frac{z_p^h (1 - a_p^h)}{E_1^h}, \quad \eta_s^h \equiv \frac{\partial E_2^h}{\partial [z_s^h (1 - a_s^h)]} \frac{z_s^h (1 - a_s^h)}{E_2^h},$$

where $a_p^h \equiv [T(z_p^h, 0, \theta) - T(0, 0, \theta)] / z_p^h$ is the participation tax rate for primary earners in household type h , and $a_s^h \equiv [T(z_p^h, z_s^h, \theta) - T(z_p^h, 0, \theta)] / z_s^h$ is the participation tax rate for secondary earners.

Because no households are observed with only the secondary earner working, government revenue can be written as

$$R = \sum_h [T(z_p^h, z_s^h, \theta) E_2^h + T(z_p^h, 0, \theta) E_1^h + T(0, 0, \theta) E_0^h],$$

which is simply the sum of the tax proceeds (net of transfers) from two-earner families (first term), one-earner families (second term), and zero-earner families (third term).

4.2 A Microsimulation Study of Reform

This section studies the effects of small tax reforms, $d\theta$, that reduce the tax burden on second-earner participation, $\partial a_s^h / \partial \theta < 0$, and are revenue-neutral, $dR/d\theta = 0$. As explained above, such reforms necessarily imply a redistribution in favor of two-earner couples at the expense of one- and zero-earner couples, and thus are associated with a trade-off between equity and efficiency. We derive theoretical measures of the equity-efficiency trade-offs associated with two specific reforms as a function of behavioral elasticities and parameters of the tax-transfer system, and apply the analytical results to our samples of married couple populations in 15 EU countries using EUROMOD.

Following Browning and Johnson (1984) and Immervoll et al. (2007), we divide the population into those who gain from the reform and those who lose from the reform. We denote by $dG \geq 0$ the aggregate welfare gain of those who gain from the reform and by $dL \leq 0$ the aggregate welfare change of those who lose from the reform. Notice that a Pareto improving reform (no losers) implies $dL = 0$, whereas a Pareto worsening reform (no gainers) implies $dG = 0$.

Due to the efficiency effects of changing distortionary taxes and transfers, the decline in welfare for those who lose from the reform (i.e., zero- and one-earner couples) is generally different from the gain in welfare for those who gain from the reform (i.e., two-earner couples). In particular, because we consider reforms designed to increase efficiency

by subsidizing second-earner participation, we would expect that the gain for two-earner couples is higher than the loss for zero- and one-earner couples. On the other hand, because two-earner couples tend to be better off than the rest of the population, policy makers may put a lower social welfare weight on the gain for two-earner couples. A critical question then becomes how to evaluate the desirability of reforms involving such inter-household utility trade-offs. The standard approach has been to postulate a social welfare function associated with certain welfare weights across different households, but the problem is that the inter-household comparisons implied by the adopted social welfare function are subjective and this limits the applicability of such an analysis as an input into the policy-making process. Following Immervoll et al. (2007), we therefore adopt a different approach, which consists in estimating critical values for the social welfare weights that would make a reform break even in terms of social welfare.

To make this precise, we define the inter-household utility trade-off Ψ in the following way:

$$\Psi = -\frac{dL}{dG}.$$

The resulting number may be interpreted as the Euro-value of the welfare loss for those who lose from the reform (zero- and one-earner couples) per additional Euro transferred to those who gain (two-earner couples). If the reform succeeds in increasing efficiency ($dL + dG > 0$), the value of Ψ is below 1, implying that it costs less than one Euro for zero-earner and one-earner couples to transfer an additional Euro to two-earner couples. However, to the extent that the social marginal welfare weight on two-earner couples

relative to other couples is below one, $\Psi < 1$ does not necessarily make the reform desirable. Generally, the lower is Ψ , the more desirable is the reform, and if $\Psi = 0$ the reform represents a Pareto-improvement.

The first reform (Reform A) reduces tax rates on secondary earners by uniformly lowering the tax burden on two-earner couples financed by uniformly increasing the tax burden on zero- and one-earner couples. The size of the extra tax on zero- and one-earner couples is determined endogenously to balance the government budget taking into account the revenue implications of behavioral responses. The reform increases second-earner participation, but has no effect on primary-earner incentives to enter the labor market as the tax increase is uniform across households with one earner and no earners. The trade-off measure for Reform A may be derived as (see Appendix C)

$$\Psi_A = \frac{1 - \sum_h e_2^h \frac{a_s^h}{1-a_s^h} \eta_s^h}{1 + \frac{e_2}{1-e_2} \sum_h e_2^h \frac{a_s^h}{1-a_s^h} \eta_s^h} < 1, \quad (6)$$

where e_2 is the share of two-earner couples in the total population of couples, and e_2^h is the share of two-earner couples that are of type h . This type of reform is always associated with an inter-household trade-off Ψ below 1: the increase in second-earner participation (at unchanged primary-earner participation) raises revenue, implying that the government can finance a welfare increase of one Euro to two-earner couples by imposing a welfare cost of less than one Euro on all other couples. It is clear from eq. (6) that the key determinants of the inter-household trade-off are the participation tax rates and participation elasticities of secondary earners, and that Ψ is decreasing in a_s^h and η_s^h as one would expect.

The second reform (Reform B) finances the tax cut on two-earner couples by taxing only one-earner couples, thereby avoiding a reduction in the welfare of zero-earner families. While reform B is associated with a better distributional profile than Reform A, the efficiency effects may be less desirable for Reform B because it increases participation tax rates on primary earners. The trade-off measure for Reform B can be expressed as (see Appendix C)

$$\Psi_B = \frac{1 - \sum_h e_2^h \frac{a_s^h}{1-a_s^h} \eta_s^h}{1 - \sum_h e_1^h \frac{a_p^h}{1-a_p^h} \eta_p^h + \frac{e_2}{e_1} \sum_h e_2^h \frac{a_s^h}{1-a_s^h} \eta_s^h}, \quad (7)$$

where e_1 is the share of one-earner households in the population, and e_1^h is the share of one-earner households that are of type h . As for the first reform, the trade-off associated with Reform B is decreasing in second-earner participation tax rates and participation elasticities. The trade-off Ψ_B additionally depends on primary-earner parameters: higher participation tax rates and higher participation elasticities for primary earners increase the trade-off. This reflects the negative efficiency effect associated with some one-earner couples dropping back to the zero-earner schedule as the tax on one-earner couples increases. Although the negative participation responses of primary earners tend to worsen the trade-off of reform B compared to reform A, there is an offsetting effect that tends to make the reform more desirable. The impact on the second-earner participation incentive is larger for reform B, because it finances the tax cuts for two-earner families entirely by higher taxes on one-earner families and therefore has a larger effect on the utility difference between two-earner and one-earner couples. Thus, it is theoretically possible that reform B improves efficiency by more than reform A, and this is more likely to occur

if the share of one-earner households e_1 is low, in which case reform B leads to a large tax increase for one-earner households.

We now turn to numerical simulations based on EUROMOD. As described, we identify the primary earner as the highest-earning member of the couple, and construct pre-tax earnings distributions for primary and secondary earners. Because the theoretical analysis is based on a discrete formulation dividing the population of couples into H earnings-groups, we have to define these subgroups in the empirical application. We divide the sample based on earnings quintiles (conditional on working) for primary and secondary earners, which yields 30 household groups (5×5 two-earner families and 5 one-earner families). For each household group, we calculate a participation tax rate using the approach described in Section 3.

We calibrate participation elasticities based on the empirical labor supply literature. There is an extensive literature on the labor force participation of married couples based on data from the United States and European countries. This literature has been surveyed by, among others, Blundell (1995) and Blundell and MaCurdy (1999). The literature finds that participation elasticities for married women (secondary earners) are substantial across a wide set of countries with values ranging from 0.5 to 1, whereas participation elasticities for prime-age males (primary earners) tend to be very small. Moreover, there is evidence that participation elasticities tend to be larger at the bottom of the earnings distribution than at the top of the earnings distribution, although some studies have found that elasticities for married women may still be substantial at the top (e.g. Eissa,

1995).

Results of the simulations are presented in Table 3. We consider four different elasticity scenarios. The first three scenarios assume that the participation elasticities are constant across earnings groups, whereas the last scenario assumes that elasticities are higher at the bottom. Average elasticities for primary and secondary earners are shown in the table for each scenario.

We start by focusing on Reform A. Recall that the inter-household trade-off associated with this reform (eq. 6) does not depend on the participation elasticity for primary earners, only the elasticity for secondary earners matters. The first scenario assumes a participation elasticity of 0.5 for secondary earners. In this scenario, many countries show a quite favorable trade-off. In Germany, one- and zero-earner couples incur a loss of just 0.14 Euros for an additional Euro distributed to two-earner couples. In Belgium, Denmark, and France, second-earner tax rates are so high that a tax cut to two-earner families creates Laffer effects and therefore a Pareto improvement. In general, the favorable trade-offs for this reform and elasticity scenario reflect the high participation tax rates on secondary earners (compared to elasticities) that we saw in Table 2. In accordance with the pattern in Table 2, Reform A is less attractive in Greece, the UK, and Spain than in Northern-Continental European countries and Scandinavia.

Not surprisingly, Reform A becomes better (worse) as the participation elasticity of secondary earners increases (declines). In the second scenario where the second-earner elasticity is set equal to 0.7, the reform is costless or nearly costless to zero- and one-earner

Table 3. Inter-household utility trade-off for reforms A and B

Country	Constant Elasticity						Elasticity Profile	
	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	$\eta_p = 0.1$	$\eta_s = 0.5$	$\eta_p = 0.1$	$\eta_s = 0.7$	$\eta_p = 0.3$	$\eta_s = 0.3$	$\eta_p = 0.1$	$\eta_s = 0.5$
	A	B	A	B	A	B	A	B
Austria	0.32	0.48	0.18	0.26	0.51	No Gainers	0.38	0.67
Belgium	No Losers	No Losers	No Losers	No Losers	No Losers	No Losers	No Losers	No Losers
Denmark	0.01	0.01	No Losers	No Losers	0.12	No Gainers	0.01	0.02
Finland	0.13	0.07	0.05	0.02	0.26	No Gainers	0.18	0.10
France	No Losers	No Losers	No Losers	No Losers	0.16	No Gainers	No Losers	No Losers
Germany	0.14	0.16	No Losers	No Losers	0.35	No Gainers	0.21	0.27
Greece	0.61	0.65	0.48	0.50	0.76	0.99	0.60	0.63
Ireland	0.27	0.31	0.06	0.07	0.52	1.73	0.30	0.39
Italy	0.33	0.32	0.14	0.13	0.55	0.80	0.34	0.34
Luxembourg	0.43	0.53	0.27	0.32	0.62	1.66	0.54	0.69
Netherlands	0.25	0.31	0.11	0.12	0.46	3.21	0.28	0.37
Portugal	0.28	0.29	0.12	0.12	0.50	0.85	0.38	0.39
Spain	0.46	0.48	0.28	0.29	0.66	0.89	0.45	0.47
Sweden	0.07	0.03	No Losers	No Losers	0.20	No Gainers	0.07	0.04
United Kingdom	0.58	0.77	0.46	0.56	0.72	No Gainers	0.65	1.03

Note: The trade-off is calculated using formula (6) in the text for reform A and formula (7) for reform B. η_p is the participation elasticity of primary earners (PE) and η_s is the participation elasticity of secondary earners (SE). Note that the primary earner elasticity does not affect the trade-off in reform A. In scenarios 1-3, the elasticities are the same for all income groups. In scenario 4, earnings responses are concentrated at the lower end of the income distribution. Specifically, η_p is 0.3 for primary earners in the lowest quintile of the PE earnings distribution (PEq1), 0.1 for PEq2 and PEq3, and 0 for PEq4 and PEq5. For secondary earners, the elasticity scenario is 1 for the lowest quintile of the SE earnings distribution (SEq1), 0.8 for SEq2, 0.5 for SEq3, 0.2 for SEq4 and 0 for SEq5. The average elasticities are listed above the results. Source: Authors' own calculations based on the EUROMOD microsimulation model.

couples in half of the countries (Belgium, Denmark, Finland, France, Germany, Ireland, and Sweden). On the other hand, in the third scenario where the second-earner elasticity is set equal to 0.3, it is only Belgium that has no losers from the reform. Nevertheless, even in this scenario, nine countries have trade-offs at or below 1/2. Scenario 4 assumes the same average elasticity as in the first scenario but with a declining profile as a function of earnings.¹⁸ The results do not change much compared to scenario 1, although there is a general tendency for the trade-off measure to increase. The reason is that the positive feedback effect on government revenue from higher participation is lower when the additional participation is generated at lower earnings levels where second-earner participation tax rates are typically lower.

The consequences of Reform B depend also on the primary-earner participation elasticity. In scenario 1, where the primary-earner elasticity is set equal to 0.1, we see that Ψ increases compared to Reform A but that the differences between the two reforms are small for all countries. Hence, the two counteracting effects on economic efficiency discussed above more or less cancel out in this elasticity scenario. When we look across the different scenarios, the effects of Reform A and B are roughly comparable except for Scenario 3 where we assume equal responsiveness for primary and secondary earners. This scenario is not realistic but highlights the importance of the relative participation elasticities when evaluating reforms of type B that affect zero- and one-earner couples

¹⁸The primary-earner elasticity is set equal to 0.3 at the lowest quintile of the primary earner income distribution (PEq1), 0.1 for PEq2 and PEq3, and 0 for PEq4 and PEq5. For secondary earners, the elasticity equals 0.8 for the lowest quintile of the secondary earner income distribution (SEq1), 0.6 for SEq2, 0.2 for SEq3, and 0 for SEq4 and SEq5.

differently. In this scenario, ten countries would experience lower efficiency by implementing reform B (i.e., $\Psi > 1$), and in seven of those countries nobody gains from the reform (Pareto worsening). The explanation is that, for most countries, primary earners face higher participation tax rates than secondary earners. This implies that, with identical elasticities, that primary-earner labor supply is more distorted than second-earner labor supply, and it is therefore suboptimal to induce additional second-earner entry at the expense of primary-earner exit.

4.3 Evaluating Reforms that Affect the Intensive Margin of Labor Supply

The reforms considered so far shift the tax burden across couples without changing marginal tax rates. Such reforms do not affect the intensive margin of labor supply (in the absence of income effects), and the assumption of fixed hours of work is therefore innocuous in the context of those reforms. But to analyze reforms associated with changes in marginal tax rates, it is necessary to extend the model to allow for both intensive and extensive responses for both spouses. Appendix D extends the model in this way, and derives the effects of a reform (Reform C) that uniformly reduces the marginal tax rate on secondary earners financed by uniformly increasing the marginal tax rate on primary earners in one-earner couples. Zero-earner couples are left unaffected. Like Reform B considered above, the reform considered here shifts participation taxes from secondary earners to primary earners, but the profile of the tax changes is different. Compared to the previous reform, changes in participation taxes are now higher at the top and lower

at the bottom.

Table 4 presents the inter-household utility trade-off implied by reform C for three different elasticity scenarios. In all three scenarios, the participation elasticities are set at our preferred levels of 0.1 for primary earners and 0.5 for secondary earners. To establish a benchmark, the first scenario assumes that hours-of-work elasticities are equal to zero for both spouses. In this case, Reform C is associated with slightly more favorable trade-offs than Reform B. The reason is that the participation tax rates of secondary earners often have an increasing profile (due to the progressivity of the tax system), whereas the participation tax rates of primary earner often display a decreasing profile (due to the impact of means-tested transfers for the first entrant). This implies that reform C (relative to Reform B) is associated with larger tax cuts to secondary earners facing the highest participation tax rates, while the tax increases tend to hit primary earners with the lowest initial participation tax rates. The second scenario sets the intensive elasticity to 0.1 for both primary and secondary earners. This generates an additional efficiency gain on the intensive margin for secondary earners, but also an efficiency loss from the intensive responses of primary earners. The total effect is that trade-offs are slightly more favorable. Scenario 3 features the same overall responsiveness on the intensive margin as Scenario 2 (the sum of the two elasticities is unchanged), but the response is now concentrated entirely on secondary earners. This makes reform C even more attractive and five countries (Belgium, Denmark, France, Germany, and Sweden) can implement the reform at no distributional cost.

Our conclusion is that the incorporation of hours-of-work responses into the analysis (assuming realistic elasticities) does not change the qualitative insights offered above and has a fairly small quantitative impact. If anything, the conclusions regarding the welfare effects of cutting taxes for secondary earners are reinforced by this generalization.

Table 4. Inter-household utility trade-off for reform C

Country	Constant Elasticity			
	Scenario 1	Scenario 2		Scenario 3
	$\epsilon_p = \epsilon_s = 0$	$\eta_p = 0.1$	$\eta_s = 0.5$	$\epsilon_p = 0 \quad \epsilon_s = 0.2$
Austria	0.42	0.34		0.18
Belgium	No Losers	No Losers		No Losers
Denmark	0.01	No Losers		No Losers
Finland	0.05	0.03		0.00
France	No Losers	No Losers		No Losers
Germany	0.11	0.00		No Losers
Greece	0.66	0.63		0.51
Ireland	0.26	0.15		0.01
Italy	0.30	0.23		0.11
Luxembourg	0.44	0.34		0.17
Netherlands	0.29	0.21		0.09
Portugal	0.18	0.11		0.03
Spain	0.48	0.45		0.35
Sweden	0.03	0.01		No Losers
United Kingdom	0.66	0.65		0.51

Note: The trade-off is calculated using formula (A-9) in Appendix D. η_p is the participation elasticity of primary earners (PE) and η_s is the participation elasticity of secondary earners (SE). Similarly, ϵ_p is the intensive elasticity of the PE and ϵ_s is the intensive elasticity of the SE. The participation elasticities are set at 0.1 for η_p and 0.5 for η_s in all scenarios. The elasticities never vary with income groups. Source: Authors' own calculations based on the EUROMOD microsimulation model.

5 Marriage Penalties in Europe

We now turn our attention to the tax-transfer implications of marriage. We present estimates of the *marriage penalty* defined as the increase in the combined tax liability net of transfers of two individuals following marriage.¹⁹ The marriage penalty has at-

¹⁹While we use the term marriage penalty throughout the paper, it would perhaps be more precise to use the label *formal cohabitation penalty*. In principle, income transfers are based on family income

tracted significant interest historically, especially in the United States where tax acts affecting married couples have often been motivated by an attempt to ‘fix’ the problem of marriage penalties. The concern about marriage penalties has been motivated by notions of fairness in the tax treatment of families (horizontal equity across married and cohabitating couples), and by the possibility that tax and transfer incentives distort the decision to marry. A number of papers have studied the effects on marriage and divorce from income taxes (e.g., Alm and Whittington, 1995a,b, 1997, 1999), welfare benefits (e.g., Hoynes, 1997a,b; Moffitt, 1998; Bitler et al., 2004), or taxes and benefits combined (Dickert-Conlin, 1999; Eissa and Hoynes, 2000b). Although these studies tend to find either modest or no effects, the implications of marriage disincentives continue to be a controversial point of contention and marriage-dependent taxes and transfers are frequently accused by conservatives of destroying the traditional two-parent family and creating higher rates of single motherhood.

Almost all existing studies of marriage penalties focus on the United States and account only for the implications of the tax system (e.g., Rosen, 1987; Feenberg and Rosen, 1995; Alm and Whittington, 1996; Dickert-Conlin and Houser, 1998; Bull et al., 1999; Alm et al., 1999; Eissa and Hoynes, 2000a). An exception to this strong US-orientation in the literature is the comparative study of marriage taxes by Pechman and Engelhardt (1990) who considered a subset of the European countries in our sample (Denmark, France, Germany, Italy, the Netherlands, Sweden, the United Kingdom).

regardless of marital status, although in practice it is difficult for the authorities to verify cohabitation when there is no marriage certificate.

While Pechman and Engelhardt considered only the tax system, we have seen in this paper that most of the jointness in redistribution schemes in Europe is driven by the welfare system suggesting that there may be important transfer-consequences to marriage. EUROMOD allows us to undertake a comparative study of marriage penalties across a large set of countries, and to incorporate fully the implication of both the tax and the transfer system.

It is helpful to start by considering some general properties of marriage penalties. Denoting by $T^i(\cdot)$ the tax function (net of transfers) that applies to individual filers and by $T^c(\cdot)$ the tax function applying to couples, the marriage penalty is defined as

$$MP \equiv T^c(z_p, z_s) - [T^i(z_p) + T^i(z_s)]. \quad (8)$$

Individual income tax treatment of couples, i.e. $T^c(z_p, z_s) = T^i(z_p) + T^i(z_s)$, is the only income tax system that does not introduce a distortion of the marriage decision, $MP = 0$. On the other hand, jointness generally implies $MP \neq 0$, and the sign of MP depends on the design of the joint schedule and on the pair of incomes z_p, z_s in a given family. If MP is negative, we say that there is a marriage subsidy. An example of a tax system giving rise to marriage subsidies (ignoring the welfare system) is a progressive and fully joint tax scheme with income splitting, so that each spouse is liable to pay taxes on half the couple's combined earnings. Formally, this is a system where $T^c(z_p, z_s) = T^i\left(\frac{z_p+z_s}{2}\right) + T^i\left(\frac{z_p+z_s}{2}\right) \equiv T^c(z_p + z_s)$. Income splitting subsidizes marriage by allowing a couple to avoid part of the progressivity of the tax system.

Family-based and means-tested transfers generally give rise to marriage penalties.

As in Section 3, let us separate the T -functions into taxes ($t^c(\cdot), t^i(\cdot)$) and benefits ($b^c(\cdot), b^i(\cdot)$). The combination of individual taxation and family-based transfers then implies $MP \equiv b^i(z_p) + b^i(z_s) - b^c(z_p + z_s)$. Then, if the $b^i(\cdot)$ and $b^c(\cdot)$ functions are the same (so that marital status is not an eligibility criterion in its own right), we have $MP > 0$ because $b'(\cdot) < 0$ as a result of means-testing. Moreover, if there is additional targeting to single parents (in which case $b^c(z) < b^i(z)$ given the presence of children), the marriage penalty is even higher. Because family-based, means-tested transfers as well as targeting to single motherhood tend to be very important at the bottom of the distribution, we would expect to find significant marriage penalties at the bottom. Moreover, these features would be particularly important in countries where welfare systems are relatively generous (such as the Nordic countries).

The marriage penalty in eq. (8) is calculated using EUROMOD by measuring the change in the combined tax liability net of transfers of a couple following a separation, holding individual earnings constant.²⁰ We consider households at ten different earnings configurations, ranging from both spouses being out of work to both spouses earning at the top decile in their respective earnings distributions, and we consider families with either two children or no children. When children are involved, we assume that each spouse takes custody of one child after the divorce.²¹ We also assume that, following the

²⁰ Although individual earnings are likely to change following a separation, it is important to keep earnings constant in order to obtain the correct tax price on marriage. Notice that, if earnings were allowed to change at separation, even a fully individual-based redistribution scheme would appear to feature marriage penalties or subsidies.

²¹ The assumption of a 1-1 split of custody is different from the usual assumptions in the (US-based) literature on marriage taxes. This literature typically assumes that either (i) the children reside with the mother (Dickert-Conlin and Houser, 1998; Eissa and Hoynes, 2000a) or (ii) custody is determined by

Table 5. Annual marriage penalties in 2007 euros for families with two children, 1998

Family Income Percentiles	AT	BE	DK	FI	FR	GE	GR	IR	IT	LU	NL	PT	SP	SW	UK
PE 0 - SE 0	5141	9770	7077	3676	14472	2387	1608	2224	0	7433	13549	0	0	10187	4696
PEp10 - SE 0	9353	10090	11640	6668	13059	5138	619	9156	-3234	12474	13375	2276	69	10495	7586
PEp10 - SEp10	7167	9712	11153	5116	7343	6573	762	10711	-1025	10624	11762	1125	837	5987	7782
PEp50 - SE 0	9353	6933	18148	12114	12096	9040	518	9999	-2307	10288	13632	2151	7	17444	13234
PEp50 - SEp10	7178	7071	11858	5611	6433	7692	615	7831	-686	4672	12343	593	849	5730	9406
PEp50 - SE p50	496	-575	9411	1055	2404	2267	57	3561	256	-3169	8814	355	1669	752	6066
PEp90 - SE 0	9353	6101	21936	14040	7308	8173	588	14318	-935	3249	12195	-187	7	17987	20780
PEp90 - SEp10	7294	6110	13457	7691	1962	7182	686	13868	-158	-2967	10713	-1621	1809	5730	17130
PEp90 - SEp50	496	-1274	10060	1055	-1192	3027	73	7023	294	-8042	6663	-1156	2821	1938	13790
PEp90 - SEp90	697	-575	6259	737	-288	955	244	5839	-422	-1027	7942	651	3316	0	2108

Note: The table shows marriage penalties on an annual basis for hypothetical families in 2007 Euros. The calculations are done for different earnings levels for the primary and secondary earner and the earnings levels refer to percentiles in the earnings distribution for primary and secondary earners, respectively. The marriage penalty is calculated as the change in a couple's combined (net-)tax liability upon separation. It is assumed that all individuals occupy rental housing. Following separation, each spouse is assumed to bear rental costs amounting to 50% of the costs in the married scenario. Further, each spouse retains custody of one child. 1998 figures are converted to 2007 euros using national indices of compensation per employee. Exchange rates are, respectively, the irrevocable euro exchange rates (12 euro countries) and 2007 average annual exchange rates (non-euro countries: Denmark, Sweden, UK). Sources: EUROMOD microsimulation model, OECD Economic Outlook 82, and Danmarks Nationalbank.

divorce, each spouse faces half the rental cost of the couple when they were married.²² The marriage penalties are shown in Table 5 for the case of two children. A corresponding table for families without children is shown in Appendix E. Marriage penalties are reported on an annual basis and in 2007 Euros.

The results reveal substantial marriage penalties in most countries, and the penalties depend primarily on the income of the lowest-earning spouse. Indeed, marriage penalties are often very high even when the primary earner is at the top decile as long as second-earner income is low. Moreover, marriage penalties are almost everywhere considerably higher when the couple has children, often more than twice as high. These patterns point to the benefit system as an important determinant of marriage penalties. In fact, in all countries, the strong targeting of transfer programs to single parents is the single most important factor contributing to marriage penalties. The tax system per se is not very important. As a result, the highest marriage penalties are found in countries that have the most generous benefit programs such as the Nordic countries, the Netherlands, France and Germany. Because of highly targeted transfers to single parents, the United Kingdom and Ireland also show substantial marriage penalties for families with children, although their social assistance programs are on the whole less generous than those of the Nordic countries. The reason is of course the high degree of targeting to lone parents

a tax minimization strategy (Rosen, 1987; Feenberg and Rosen, 1995). Because the second assumption implies that typically the higher-earnings spouse takes custody of all the children, whereas the first assumption implies that typically the lower-earnings spouse gets the kids, our assumption of an equal split lies in between these two extremes.

²²Our approach to the calculation of marriage penalties is closely related to the so-called “Resource Pooling Approach” (see Bull et al., 1999). Our calculations do not include unearned income and therefore capture only the marriage penalty arising from the tax-transfer treatment of earned income.

in the Anglo-Saxon system.

There are some exceptions to this general pattern of high marriage penalties. Italy offers non-trivial marriage subsidies resulting from family benefits available only to married couples with children. The tax-transfer system in Greece is virtually neutral with respect to marriage for couples without children but does feature minor penalties for couples with children. This is the result of an individual income tax combined with fairly small social assistance benefits that are available only to single parents. Spain tends to subsidize marriage for couples without children but penalize it for couples with children. In France, marriage subsidies are considerable for higher-income families.²³

6 Conclusion

The standard Mirrleesian theory of optimal income taxation assumes that all tax payers and transfer recipients live in single-person households. In reality, most individuals live in families, and the tax-transfer rules applying to a married person are often different from the rules applying to single individuals. A number of recent papers have attempted to generalize the theory of optimal income taxation to explicitly deal with couples. Instead of characterizing the optimal tax-transfer treatment of families, this paper characterizes the actual tax-transfer treatment of couples and identifies welfare-improving reforms for

²³An additional important factor determining marriage penalties are housing benefits. In results not reported here, we have calculated marriage penalties under alternative assumptions about housing costs following separation. For example, if rental costs for each spouse are at the same level as the combined rental costs of the couple (reflecting economies of scale in two-person households), the size of marriage penalties are considerably affected in some countries (in particular, the UK, Ireland, Germany, Austria, Denmark, Finland, and the Netherlands).

15 European countries.

We have considered three aspects of a tax-transfer system for couples: the form of jointness, the distortion of second-earner labor supply, and the size of marriage penalties. A general insight from the analysis is that the transfer system is a crucial element in understanding and evaluating redistribution schemes affecting married couples. For example, it is the presence of family-based and means-tested transfer programs that explains the observation in many countries of negative jointness, i.e. a negative relationship between the effective tax rate on one person and the earnings of the spouse. Interestingly, negative jointness is in accordance with prescriptions from the recent optimal tax literature (Kleven, Kreiner and Saez, 2007, 2008). At the same time, family-based transfers tend to create substantial marriage penalties at the bottom of the distribution, which raise issues pertaining to fairness and to some extent efficiency.

Our analysis of tax-transfer distortions at the extensive margin of labor supply suggests that the effective taxation of secondary earners relative to primary earners is too high given the empirical evidence on participation elasticities. Simple revenue-neutral reforms that shift some of the tax burden from two-earner couples to one-earner and/or zero-earner couples could reduce the distortion of second-earner labor supply and may generate large welfare gains. In fact, for some countries, a tax cut on secondary earners may realistically pay for itself and give rise to a Pareto improvement. For countries where Laffer effects are not present, a tax cut for two-earner families does require a higher tax on other couples, but the required tax increase tends to be reasonably small. In a ma-

majority of countries, it is possible to transfer 1 euro to two-earner couples by taking away less than $\frac{1}{2}$ a euro from one- and no-earner couples. If the reform is financed by raising taxes for one-earner couples only, their burden will be somewhat higher but this avoids making the poorest families worse off.

Appendix A: Marginal Tax Rates and Jointness

This section presents evidence on marginal tax rates that corresponds to the information on participation tax rates found in Table 2 and Figure 1.

Table A1. Average marginal tax rates

Country	Primary earners	Secondary earners	PE tax / SE tax
Austria	0.52	0.44	1.19
Belgium	0.65	0.60	1.08
Denmark	0.58	0.54	1.07
Finland	0.54	0.47	1.16
France	0.53	0.53	0.99
Germany	0.54	0.59	0.92
Greece	0.37	0.28	1.32
Ireland	0.43	0.43	1.01
Italy	0.52	0.48	1.07
Luxembourg	0.41	0.44	0.95
Netherlands	0.52	0.47	1.12
Portugal	0.44	0.42	1.03
Spain	0.37	0.32	1.17
Sweden	0.59	0.54	1.09
United Kingdom	0.37	0.31	1.21

Note: Average effective marginal tax rates for primary earners and secondary earners conditional on working. Source: EUROMOD Microsimulation Model.

Figure A1. Marginal tax rates of secondary earners (in percent) at selected earnings levels as a function of primary earnings (vingtiles) for couples with two children


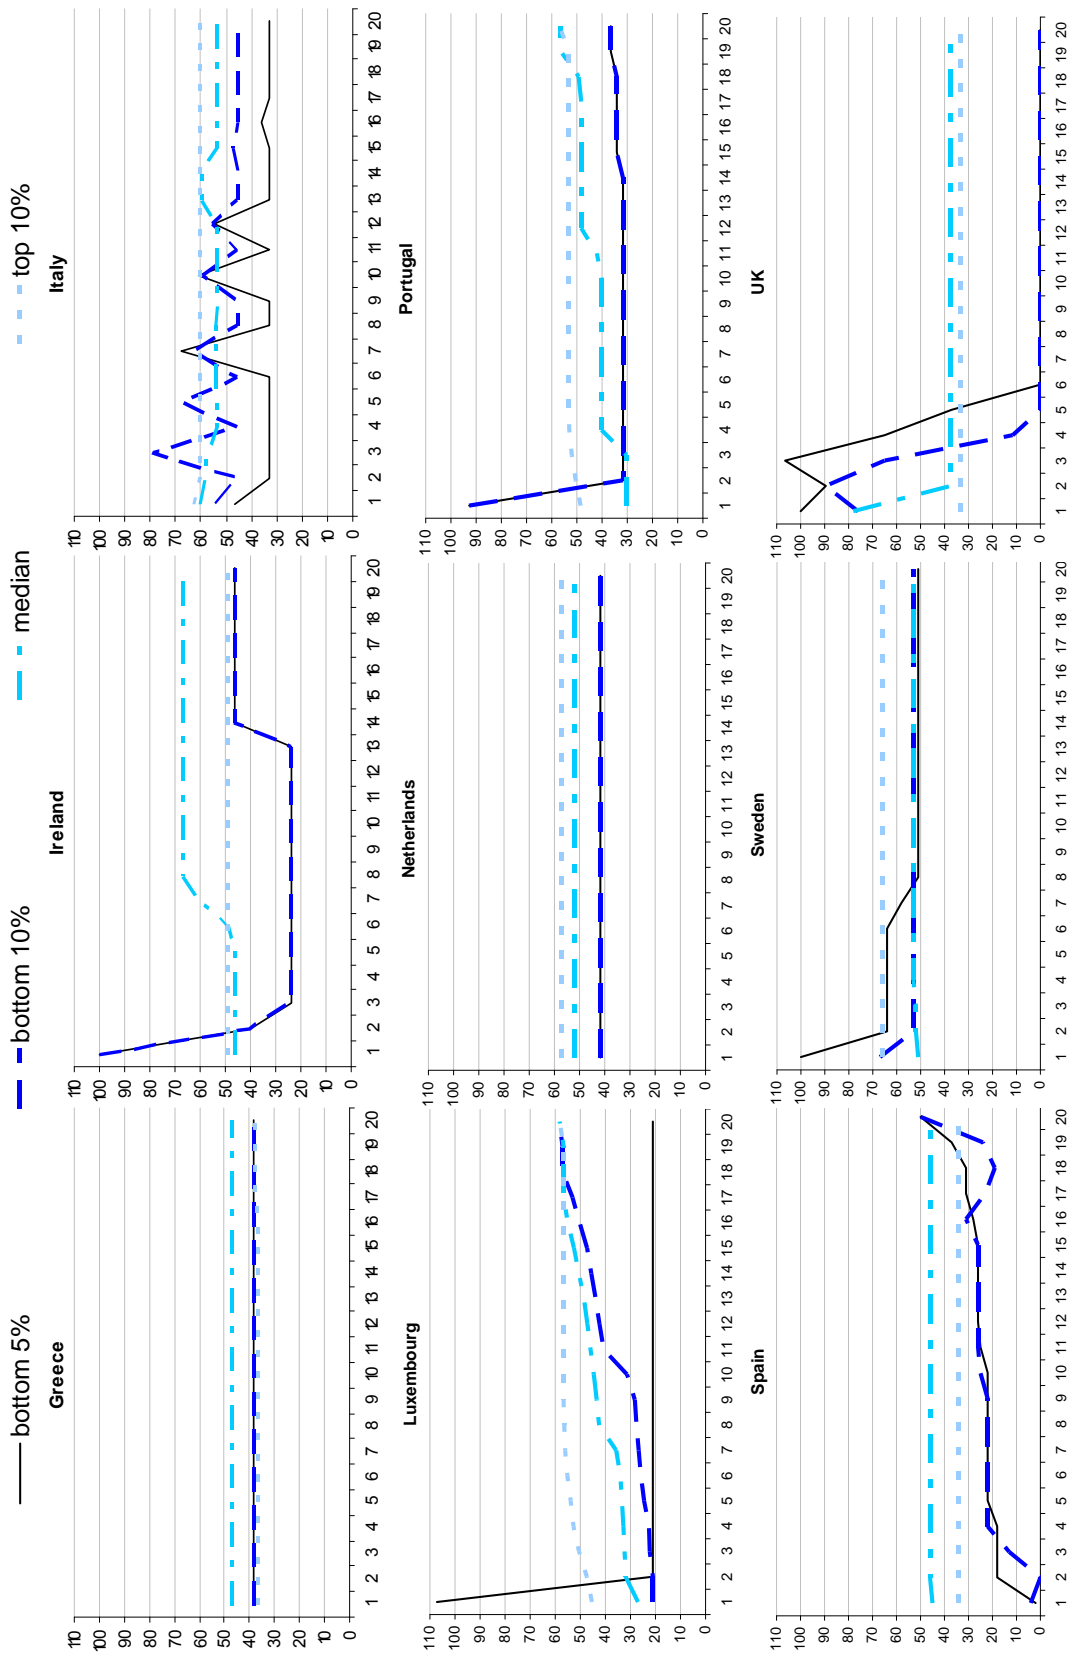


Figure continues on the next page

Figure A1 cont.



Note: Each figure shows the marginal tax rate of the secondary earner (SE) in percent as a function of the earnings of the primary earner (PE), depicted by vingtile group in the PE earnings distribution. The graphs are shown for four different SE earnings levels: the lowest vingtile (SEP5), the lowest decile (SEP10), the median (SEP50), and the top decile (SEP90) of the SE earnings distribution. Source: EUROMOD microsimulation model.

Appendix B: Figure A2. Participation tax rates of secondary earners (in percent) at selected earnings levels as a function of primary earnings (vingtiles) for couples with no children

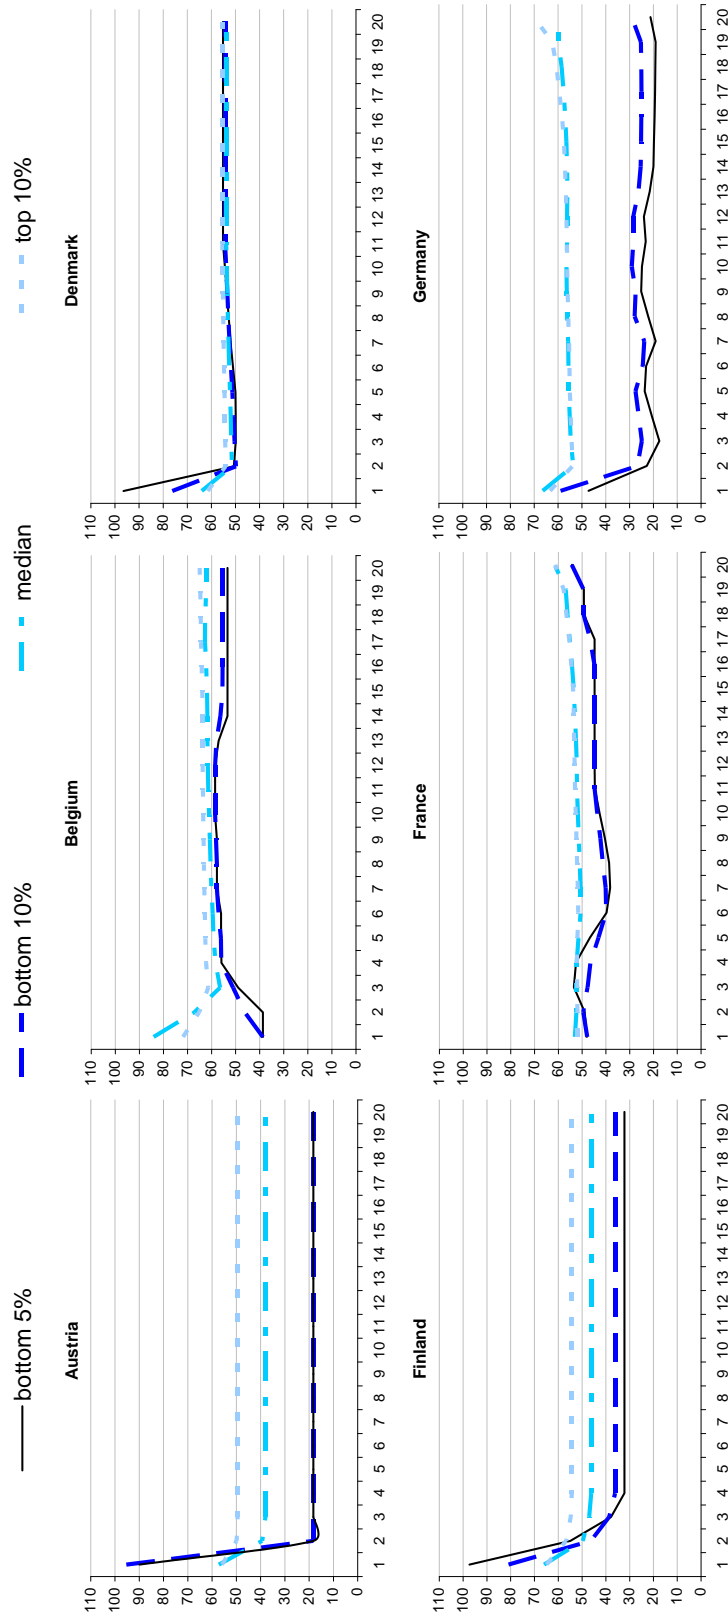



Figure continues on the next page

Figure A2 cont.



Note: Each figure shows the participation tax rate of the SE in percent as a function of the earnings of the PE, depicted by vingtile group in the PE earnings distribution. The graphs are shown for four different SE earnings levels: the lowest vingtile (SEp5), the lowest decile (SEp10), the median (SEp50), and the top decile (SEp90) of the SE earnings distribution. Source: EUROMOD microsimulation model.

Appendix C: Derivation of eqs (6) and (7)

We derive the inter-household utility trade-off under the assumption that there are no households with only the secondary earner working either before or after a reform. To ensure that this is consistent with household optimization we must restrict the distribution of the fixed costs of work for secondary earners. In terms of Figure 2, we must make sure that a marginal reform does not induce any families to position themselves in area \emptyset . We denote by $V_h(\cdot)$ the indirect family utility function, which depends on the work status of the two spouses. The conditions we will impose on the distribution of fixed costs of work for the secondary earner amount to saying that, following a marginal reform, the indirect utility is greater for *all* families if they are in area 0 of Figure 2 than if they are in area \emptyset . A sufficient condition makes sure that no couples have a high fixed cost of work for the primary earner q_p and at the same time a relatively low q_s for the secondary earner.

Let $\bar{q}_s^h(0) \equiv V_h(0, 1) - V(0, 0)$ be the gain from secondary earner entry for household type h when the primary earner is *not* working, and let $\bar{q}_p^h(1) \equiv V_h(1, 1) - V_h(0, 1)$ be the gain from primary earner entry when the secondary earner is *already* working. Further, let $\underline{q}_p^h \equiv \min\{\bar{q}_p^h(1), \bar{q}_p^h\}$ where \bar{q}_p^h is determined by (4). We will assume a lower bound on the secondary earner fixed costs of work depending on the primary earner's fixed costs of work, $\underline{q}_s^h(q_p)$, such that $P_h(\underline{q}_s^h(q_p) | q_p) \equiv 0$. The lower bound assumption is

$$\underline{q}_s^h(q_p) > q_s^h(0) \text{ for } q_p > \underline{q}_p^h. \quad (\text{A-1})$$

The reason for the two different thresholds for the primary earner is that we must consider both the potential movement from area 2 to area \emptyset and from area 0 to area \emptyset in Figure 2.

With this assumption, there will be no households with only the secondary earner working either before or after marginal reforms. Government revenue can then be written as

$$\begin{aligned}
 R &= \sum_h N_h \left[\int_0^{\bar{q}_p^h} \int_{\underline{q}_s^h(q_p)}^{\bar{q}_s^h} T(z_p^h, z_s^h, \theta) p_h(q_s | q_p) f_h(q_p) dq_s dq_p \right. \\
 &\quad + \int_0^{\bar{q}_p^h} \int_{\bar{q}_s^h}^{\infty} T(z_p^h, 0, \theta) p_h(q_s | q_p) p_h(q_s | q_p) f_h(q_p) dq_s dq_p \\
 &\quad \left. + \int_{\bar{q}_p^h}^{\infty} \int_{\underline{q}_s^h(q_p)}^{\infty} T(0, 0, \theta) p_h(q_s | q_p) f_h(q_p) dq_s dq_p \right] \\
 &= \sum_h [T(z_p^h, z_s^h, \theta) E_2^h + T(z_p^h, 0, \theta) E_1^h + T(0, 0, \theta) E_0^h], \quad (\text{A-2})
 \end{aligned}$$

where E_2^h is the number of two-earner households of type h , or equivalently, the number of working secondary earners of type h , E_1^h is the number of one-earner households of type h , and E_0^h is the number of type h households without any labor force attachment.

The effect of a small reform, $d\theta$, on government revenue is given by

$$\begin{aligned}
 \frac{dR}{d\theta} &= \sum_h \left[\frac{\partial T_h(1, 1)}{\partial \theta} E_2^h + \frac{\partial T_h(1, 0)}{\partial \theta} E_1^h + \frac{\partial T_h(0, 0)}{\partial \theta} E_0^h \right. \\
 &\quad \left. + T_h(1, 1) \frac{dE_2^h}{d\theta} + T_h(1, 0) \frac{dE_1^h}{d\theta} + T_h(0, 0) \frac{dE_0^h}{d\theta} \right] \\
 &= \sum_h \left[\frac{\partial T_h(1, 1)}{\partial \theta} E_2^h + \frac{\partial T_h(1, 0)}{\partial \theta} E_1^h + \frac{\partial T_h(0, 0)}{\partial \theta} E_0^h \right. \\
 &\quad \left. + [T_h(1, 1) - T_h(1, 0)] \frac{dE_2^h}{d\theta} + [T_h(1, 0) - T_h(0, 0)] \left(\frac{dE_1^h}{d\theta} + \frac{dE_2^h}{d\theta} \right) \right], \quad (\text{A-3})
 \end{aligned}$$

because

$$\frac{dE_0^h}{d\theta} = - \left(\frac{dE_1^h}{d\theta} + \frac{dE_2^h}{d\theta} \right).$$

The first three terms in (A-3) reflect mechanical effects while the last two terms capture the effects of the behavioral responses. The mechanical revenue effects are simply the direct effects on tax revenue with unchanged behavior. The behavioral responses constitute new entry of secondary earners (term 4) as well as new entry by primary earners (term 5).

The employment effects of the reform can be expressed using the participation elasticities from the main text. The change in total employment among couples is

$$\frac{dE_1^h}{d\theta} + \frac{dE_2^h}{d\theta} = \frac{\partial E_1^h}{\partial \bar{q}_p^h} \frac{d\bar{q}_p^h}{d\theta} = - \frac{\partial E_1^h}{\partial z_p^h (1 - a_p^h)} z_p^h \frac{\partial a_p^h}{\partial \theta} = - \frac{1}{1 - a_p^h} \frac{\partial a_p^h}{\partial \theta} \eta_p^h E_1^h, \quad (\text{A-4})$$

where we have used

$$\frac{\partial a_p^h}{\partial \theta} = \left(\frac{\partial T_h(1, 0)}{\partial \theta} - \frac{\partial T_h(0, 0)}{\partial \theta} \right) / z_p^h(0). \quad (\text{A-5})$$

The change in secondary employment is

$$\frac{dE_2^h}{d\theta} = \frac{\partial E_2^h}{\partial \bar{q}_s^h} \frac{d\bar{q}_s^h}{d\theta} = - \frac{\partial E_2^h}{\partial z_s^h (1 - a_s^h)} z_s^h \frac{\partial a_s^h}{\partial \theta} = - \frac{1}{1 - a_s^h} \frac{\partial a_s^h}{\partial \theta} \eta_s^h E_2^h, \quad (\text{A-6})$$

where use is made of

$$\frac{\partial a_s^h}{\partial \theta} = \left(\frac{\partial T_h(1, 1)}{\partial \theta} - \frac{\partial T_h(1, 0)}{\partial \theta} \right) / z_s^h \quad (\text{A-7})$$

Using (A-4) and (A-6) we can rewrite (A-3) as

$$\begin{aligned} \frac{dR}{d\theta} &= \sum_h \left[\frac{\partial T_h(1, 1)}{\partial \theta} E_2^h + \frac{\partial T_h(1, 0)}{\partial \theta} E_1^h + \frac{\partial T_h(0, 0)}{\partial \theta} E_0^h \right] \\ &\quad - \sum_h \left[\frac{a_p^h}{1 - a_p^h} \frac{\partial a_p^h}{\partial \theta} \eta_p^h z_p^h E_1^h + \frac{a_s^h}{1 - a_s^h} \frac{\partial a_s^h}{\partial \theta} \eta_s^h z_s^h E_2^h \right]. \end{aligned}$$

Reform A. The first reform has

$$\frac{\partial T_h(1,0)}{\partial \theta} = \frac{\partial T_h(0,0)}{\partial \theta} = \frac{\partial T(1,0)}{\partial \theta} > 0, \quad \frac{\partial T_h(1,1)}{\partial \theta} = \frac{\partial T(1,1)}{\partial \theta} < 0 \quad \forall h,$$

where the tax increase to one- and no-earner families is determined endogenously. Because the reform is purely redistributive, government revenue must remain unchanged, $\frac{dR}{d\theta} = 0$, implying

$$-\frac{\partial T(1,0)}{\partial \theta} = \frac{\frac{\partial T(1,1)}{\partial \theta} E_2 \left(1 - \sum_h \frac{E_2^h}{E_2} \frac{a_s^h}{1-a_s^h} \eta_s^h\right)}{N - E_2 \sum_h \frac{E_2^h}{E_2} \left(1 - \frac{a_s^h}{1-a_s^h} \eta_s^h\right)},$$

where $E_2 = \sum_h E_2^h$ is the total number of two-earner households in the economy and where we have used (A-7). From the envelope theorem and the marginal nature of the reform, monetary gains and losses are simply the direct changes in tax liabilities. Because all two-earner couples gain and all zero- and one-earner couples lose, we have

$$\Psi_A = -\frac{-\sum_h \frac{\partial T(1,0)}{\partial \theta} (N_h - E_2^h)}{-\sum_h \frac{\partial T(1,1)}{\partial \theta} E_2^h} = \frac{-\frac{\partial T(1,0)}{\partial \theta} (N - E_2)}{\frac{\partial T(1,1)}{\partial \theta} E_2} = \frac{(N - E_2) \left(1 - \sum_h \frac{E_2^h}{E_2} \frac{a_s^h}{1-a_s^h} \eta_s^h\right)}{N - E_2 \sum_h \frac{E_2^h}{E_2} \left(1 - \frac{a_s^h}{1-a_s^h} \eta_s^h\right)}.$$

By inserting $e_2 \equiv E_2/N$ and $e_2^h \equiv E_2^h/E_2$, we obtain expression (6).

Reform B. The second reform implies

$$\frac{\partial T_h(0,0)}{\partial \theta} = 0, \quad \frac{\partial T_h(1,0)}{\partial \theta} = \frac{\partial T(1,0)}{\partial \theta} > 0, \quad \frac{\partial T_h(1,1)}{\partial \theta} = \frac{\partial T(1,1)}{\partial \theta} < 0 \quad \forall h,$$

again with the tax increase for two-earner couples exogenously given. We find

$$-\frac{\partial T(1,0)}{\partial \theta} = \frac{\frac{\partial T(1,1)}{\partial \theta} E_2 \left(1 - \sum_h \frac{E_2^h}{E_2} \frac{a_s^h}{1-a_s^h} \eta_s^h\right)}{E_1 \left(1 - \sum_h \frac{E_1^h}{E_1} \frac{a_p^h}{1-a_p^h} \eta_p^h\right) + \sum_h E_2^h \frac{a_s^h}{1-a_s^h} \eta_s^h},$$

where $E_1 = \sum_h E_1^h$ is the number of one-earner households and where we have used eqs (A-5) and (A-7) as well as the fact that only primary earners in one-earner couples respond

to the reform. The trade-off between equity and efficiency equals $\Psi_B = -\frac{-\frac{\partial T(1,1)}{\partial \theta} E_2}{-\frac{\partial T(1,0)}{\partial \theta} E_1}$. By inserting the derivatives from above and the definitions $e_1 \equiv E_1/N$, $e_1^h = E_1^h/E_1$, $e_2 \equiv E_2/N$ and $e_2^h \equiv E_2^h/E_2$, we obtain expression (7).

Appendix D: The intensive-extensive model

We introduce intensive responses by allowing individuals to choose working hours subject to the costs of working time given by $v_p(z_p/n_p)$ for primary earners and $v_s(z_s/n_s)$ for secondary earners, where z/n is working time for an individual with earnings z and innate ability n . The household utility function is now given by

$$u(c, z_p, z_s) = c - n_p v_p\left(\frac{z_p}{n_p}\right) - n_s v_s\left(\frac{z_s}{n_s}\right) - q_p \cdot 1(z_p > 0) - q_s \cdot 1(z_s > 0).$$

Conditional on working, the primary earner chooses working hours according to

$$1 - m_p^h(l) = v_p' \left(\frac{z_p^h(l)}{n_p^h} \right) \text{ for } l = 0, 1,$$

where $m_p^h(\cdot) \equiv T_p'$ is the marginal tax rate faced by the primary earner, which may depend on the work status of the secondary earner. Thus, $l = 1$ denotes a working spouse and $l = 0$ represents a non-working secondary earner. Similarly, the number of working hours for the secondary earners conditional on participation satisfies

$$1 - m_s^h = v_s' \left(\frac{z_s^h}{n_s^h} \right),$$

where $m_s^h \equiv T_s'$ denotes the marginal tax rate for the secondary earner (which does not depend on the work status of the primary earner because a working secondary earner is

always married to a working primary earner). Household behavior along the intensive margin is captured by the intensive elasticities

$$\varepsilon_p^h(l) = \frac{1 - m_p^h(l)}{z_p^h(l)} \frac{\partial z_p^h(l)}{\partial [1 - m_p^h(l)]} \text{ for } l = 0, 1, \quad \varepsilon_s^h = \frac{1 - m_s^h}{z_s^h} \frac{\partial z_s^h}{\partial (1 - m_s^h)},$$

which give the change in earnings in response to a change in the net-of-tax rate rates for the primary and the secondary earner, respectively. Since the marginal tax rates of an individual may now depend on spousal income, so may the choice of earnings. In particular, the earnings of the primary earner are likely to change when the secondary earner enters the labor market. Behavior along the extensive margin is governed by the exact same logic as in the simpler model: the primary earner in household h enters whenever entry increases the family's utility, i.e., when $V_h(1, 0) - V(0, 0) \geq q_p$. Similarly, the secondary earner in household h enters whenever $V_h(1, 1) - V_h(1, 0) \geq q_s$.²⁴ Compared to the simpler model, the correct definition of the secondary earner participation tax rate now includes the tax implications of the change in primary earnings, i.e., $a_s^h = T(z_p^h(1), z_s^h, \theta) - T(z_p^h(0), 0, \theta) / z_s^h$.²⁵

The definition of government revenue R is unchanged. The change in R as a result of

²⁴Assumption (A-1) is again sufficient to solve the double screening problem.

²⁵In the empirical simulations, we are forced to assume that primary earnings remain unchanged when the secondary earner enters.

the reform, $d\theta$, is

$$\begin{aligned} \frac{dR}{d\theta} = & \sum_h \left[\frac{\partial T_h [1, 1]}{\partial \theta} E_2^h + \frac{\partial T_h [1, 0]}{\partial \theta} E_1^h + \frac{\partial T_h (0, 0)}{\partial \theta} E_0^h \right. \\ & + [T_h (1, 1) - T_h (1, 0)] \frac{dE_2^h}{d\theta} + [T_h (1, 0) - T_h (0, 0)] \left(\frac{dE_1^h}{d\theta} + \frac{dE_2^h}{d\theta} \right) \\ & \left. + m_p^h (1) \frac{dz_p^h (1)}{d\theta} E_2^h + m_s^h \frac{dz_s^h}{d\theta} E_2^h + m_p^h (0) \frac{dz_p^h (0)}{d\theta} E_1^h \right]. \end{aligned} \quad (\text{A-8})$$

As before, the employment effects can be rewritten using elasticities to find

$$\begin{aligned} \frac{dR}{d\theta} = & \sum_h \left[\frac{\partial T_h (1, 1)}{\partial \theta} E_2^h + \frac{\partial T_h (1, 0)}{\partial \theta} E_1^h + \frac{\partial T_h (0, 0)}{\partial \theta} E_0^h \right. \\ & - \frac{a_p^h}{1 - a_p^h} \frac{\partial a_p^h}{\partial \theta} \eta_p^h z_p^h (0) E_1^h - \frac{a_s^h}{1 - a_s^h} \frac{\partial a_s^h}{\partial \theta} \eta_s^h z_s^h E_2^h \\ & - \frac{m_p^h (1)}{1 - m_p^h (1)} \frac{\partial m_p^h (1)}{\partial \theta} \varepsilon_p^h (1) z_p^h (1) E_2^h \\ & - \frac{m_s^h}{1 - m_s^h} \frac{\partial m_s^h}{\partial \theta} \varepsilon_s^h z_s^h E_2^h \\ & \left. - \frac{m_p^h (0)}{1 - m_p^h (0)} \frac{\partial m_p^h (0)}{\partial \theta} \varepsilon_p^h (0) z_p^h (0) E_1^h \right]. \end{aligned}$$

Reform C. The details of the reform are

$$\begin{aligned} \frac{\partial m_p^h (1)}{\partial \theta} = \frac{\partial T_h (0, 0)}{\partial \theta} = 0, \quad \frac{\partial m_p^h (0)}{\partial \theta} = \tau \implies \frac{\partial T_h (1, 0)}{\partial \theta} = \tau z_p^h (0), \\ \frac{\partial m_s^h}{\partial \theta} = -t \implies \frac{\partial T_h (1, 1)}{\partial \theta} = -t z_s^h \quad \forall h, \end{aligned}$$

where $t > 0$ is exogenous while τ is endogenously determined by government budget neutrality, $dR = 0$. This implies

$$\tau = \frac{t \sum_h z_s^h E_2^h \left(1 - \frac{a_s^h}{1 - a_s^h} \eta_s^h - \frac{m_s^h}{1 - m_s^h} \varepsilon_s^h \right)}{\sum_h z_p^h (0) E_1^h \left(1 - \frac{a_p^h}{1 - a_p^h} \eta_p^h + \frac{E_2^h}{E_1^h} \frac{a_s^h}{1 - a_s^h} \eta_s^h - \frac{m_p^h (0)}{1 - m_p^h (0)} \varepsilon_p^h (0) \right)}.$$

Using this expression, the trade-off between equity and efficiency becomes

$$\begin{aligned} \Psi_C &= -\frac{-\sum_h \frac{\partial T_h(1,0)}{\partial \theta} E_p^h}{-\sum_h \frac{\partial T_h(1,1)}{\partial \theta} E_s^h} = \frac{\tau \sum_h z_p^h(0) E_1^h}{t \sum_h z_s^h E_2^h} \\ &= \frac{1 - \sum_h s_s^h \left(\frac{a_s^h}{1-a_s^h} \eta_s^h + \frac{m_s^h}{1-m_s^h} \varepsilon_s^h \right)}{1 - \sum_h s_p^h \left(\frac{a_p^h}{1-a_p^h} \eta_p^h - \frac{e_2^h e_1}{e_1^h e_1} \frac{a_s^h}{1-a_s^h} \eta_s^h + \frac{m_p^h(0)}{1-m_p^h(0)} \varepsilon_p^h(0) \right)}, \end{aligned} \quad (\text{A-9})$$

where $s_p^h \equiv z_p^h(0) E_p^h / (\sum_h z_p^h(0) E_p^h)$ is the share of all earnings in one-earner families that accrues to households of type h , and $s_s^h \equiv z_s^h E_s^h / \sum_h z_s^h E_s^h$ is the share of all secondary earnings accruing to household h .

Appendix E: Marriage Penalties for Couples Without Children

Table A2. Annual marriage penalties in 2007 euros for families with no children, 1998

Family Income Percentiles	AT	BE	DK	FI	FR	GE	GR	IR	IT	LU	NL	PT	SP	SW	UK
PE 0 - SE 0	4905	5217	-2990	2066	5907	760	0	2224	0	8110	2443	0	0	10384	2418
PEp10 - SE 0	7510	5319	4966	7263	7306	4008	-24	7069	-1527	13281	6312	1652	-421	14126	4994
PEp10 - SEp10	4875	5684	199	2214	2440	2587	-24	4471	-437	8298	5385	94	0	2251	3336
PEp50 - SE 0	7780	5861	8184	8340	6915	3273	-145	5078	-957	9089	6190	1241	-969	14508	8075
PEp50 - SEp10	5228	6741	264	2224	1485	1913	-145	1235	-456	3201	4886	-317	-127	2251	3587
PEp50 - SE p50	0	1655	0	0	-170	-149	11	-1507	-625	-3827	961	142	0	0	-609
PEp90 - SE 0	8020	6027	8985	9621	1092	-356	0	5922	-1619	3121	8937	-999	-2122	14508	13411
PEp90 - SEp10	5969	6779	1031	3223	-3263	-1468	0	2698	-1207	-3539	7560	-2434	-320	2251	9049
PEp90 - SEp50	0	1954	1464	0	-3262	-5517	106	-1212	-1377	-7801	0	-1169	0	0	3148
PEp90 - SEp90	0	3077	0	0	-1251	-1682	287	0	-1677	-3662	0	638	0	0	-610

Note: The table shows marriage penalties on an annual basis for hypothetical families in 2007 Euros. The calculations are done for different earnings levels for the primary and secondary earner and the earnings levels refer to percentiles in the earnings distribution for primary and secondary earners, respectively. The marriage penalty is calculated as the change in a couple's combined (net-)tax liability upon separation. It is assumed that all individuals occupy rental housing. Following separation, each spouse is assumed to bear rental costs amounting to 50% of the costs in the married scenario. Further, each spouse retains custody of one child. 1998 figures are converted to 2007 euros using national indices of compensation per employee. Exchange rates are, respectively, the irrevocable euro exchange rates (12 euro countries) and 2007 average annual exchange rates (non-euro countries: Denmark, Sweden, UK). Sources: EUROMOD microsimulation model, OECD Economic Outlook 82, and Danmarks Nationalbank.

Table A3: Summary of taxes on workers

	Income Tax ⁷					SSC (employee)				SSC (employer) and Payroll Tax				features reducing METR and PTR	features increasing METR and PTR
	lowest/highest tax band limit ^{1,2}	lowest/highest rate [%]	main tax credit ¹	tax unit	family-related tax provisions	thres-hold ¹	rate [%]	ceiling ¹	tax deductible	thres-hold ¹	rate [%]	ceiling ¹	taxable		
Austria	17/231	21/50 4 rates	5	individual	deduction for single earners; tax credits for lone parents	15	18.8	193	yes	15	21.3 4.5	193	no		2nd earner: phase-out of single-earner credit
Belgium	24/318	25/55 7 rates	-	individual	parts of taxable income transferrable to spouse; additional tfa for children and lone parents	-	11.9	-	yes	-	45.4	-	no		2nd earner: phase-out of amount transferrable from higher-earning spouse
Denmark	12/100	40/59 ⁴ 3 rates	-	individual	unused deductions transferrable to spouse	-	9 + flat amount	-	yes	-	2 + flat amount	-	no		
Finland	35/223	24/56 ⁴ 6 rates	-	individual		- 58	7.6 0.5	-	yes	-	24.5	-	no	earned income tax allowance (20% of earnings above 11)	phase-out of earned income tax and basic allowances (starting at earnings > 31 and 76 respectively)
France	30/336	11/54 6 rates	-	family		- - - 136	0.9 9.6 2.8 ⁴ 3.6 7.6	- 136 409 545 -	yes yes yes yes partly	- - - 136	19.8 13.4 4.1 5.3	- 136 409 545	no	reduced employer contributions rates for wages < 130% MW.	
Germany	30 133 252	27.3 37.2 55.7 ⁵	-	married couple (individual optional)	choice of tfa or child benefit	15	7.7 13.4	156 208	yes	15	7.7 13.4	156 208	no		
Greece	56/478	5/45 5 rates	max. 15% of accepted hshld. expenditure	individual	0.9-1.8 non-refundable tax credit per child	-	15.9	200; none for new jobs	yes	-	28.2	200; none for new jobs	no		
Ireland	25/80	24/46 2 rates	-	married couple (individual optional)	-	41 86	4.5 2.3	193 -	no	- 112 -	8.5 12.0 4.0	12 231 231	no	no tax if income below 33 (higher limit if children)	tax-free status phased out above the 33 limit.
Italy	0/118	19/46 5 rates	up to 6	individual	up to 2 tax credit for each dependent family member	- 56	9.0 ⁴ 1	-	yes	-	33.0 ⁴	-	no		phase-out of main tax credit; 2nd earner: phase-out of tax credit for dependent spouse;
Luxembourg	25/250	6/47 17 rates	-	married couple	deductions for lone parents and care expenditure; 3 tax credit per child	-	13.1	259	yes	-	14.6 ⁴	259	no	2nd earner: additional tax deduction if both spouses work	
Netherlands	20/212	36 ⁸ /60 3 rates	-	individual	additional tfa for lone parents	54 -	5.3 1.7	156 105	yes no	54 -	6.4 5.6 7.9 ⁴	156 105 156	no yes no		
Portugal	0/490	5 ⁹ /40 5 rates	3	married couple	additional 1.5 tax credit per child	-	11	-	no	-	23.8	-	no		
Spain	22/492	20/56 8 rates	3	family (individual optional)	up to 2 tax credit per child (additional amounts in some regions)	46	30.8	177	yes	46	6.4	177	no	earning < 55 are tax-exempt	"spike" in METR once above exemption limit; phase-out of main tax credit adds 5 pct. points to METR
Sweden	4/92	30 ⁴ 55 ⁴	-	individual	-	3	6.95	121	yes	0.4	33	-	no		
United Kingdom	29/220	20/40 3 rates	-	individual	2 tax credit for married couples; 13 tax deduction for lone parents	23	8.4 to 10	177	no	23 ¹⁰ 40 56 76	3 5 7 10	-	no		

Appreviations: SSC = social security contribution; METR = marginal effective tax rate; PTR = participation tax rate; tfa = tax free allowance; MW = statutory minimum wage.

Explanatory Notes: Reference year is 1998. Except where noted, all information is for private-sector employees with no other income and not claiming itemised expenses. Income taxes include local and regional taxes where applicable. Multiple lines of SSC entries are shown where payment schedules differ for the different programs (e.g. for pensions, health, unemployment, etc.). Further information, and data for later years, can be found in the EUROMOD country reports (www.iser.essex.ac.uk/msu/emod/documentation/countries/), the OECD Tax Database (www.oecd.org/ctp/taxdatabase) and the OECD series *Benefits and Wages* (www.oecd.org/els/social/workincentives).

¹ in % of median gross earnings of primary earners (not including employer social security contributions)

² after adding any standard tax free allowances, deductions or exemptions available to single employees

³ insurance is voluntary

⁴ averages: rates differ between municipalities and/or employers

⁵ including "Solidarity Surplus Tax" for German unification. MTR increases linearly inbetween lower and middle; and middle and top tax band limits.

⁶ West Germany

⁷ including regional income taxes where applicable

⁸ including pension contributions (same tax base as income tax)

⁹ effective rate taking into account the allowance of 70% of the tax base for low incomes

¹⁰ all earnings are subject to the applicable rate once they exceed these threshold levels

Table A3: Summary of taxes on workers

	Social Assistance				Housing Benefit ²		Family Benefits ³		Employment-conditional Benefits		
	max. amount ¹	disregard ¹	withdrawal rate	taxable	max. amount ¹	withdrawal rate	amount ¹	withdrawal rate	amount ¹	work/income conditions	withdrawal rate
Austria	52	-	100%	IT: no SSC: no	-	-	5-7 per child	none (universal payment)	-	-	-
Belgium	53	10 for each working adult	100%	IT: no SSC: no	-	-	4-13 per child	none (universal payment)	-	-	-
Denmark	90 + housing allowance	9 for each working adult	100%	IT: yes SSC: no	11	75%	3-4 per child; higher for lone parents	none (universal payment)	-	-	-
Finland	58 +reasonable housing cost	-	100%	IT: no SSC: no	34	80%	5-9 per child; plus 2 per child for lone parents; plus day-care subsidy	none (universal payment)	-	-	-
France	49	-	100%	IT: no SSC: no	20	34%	main benefit: 7 to 12 for second & further children; special benefits for young children	main benefit: 100% once income > 174-261	-	-	-
Germany	47 ⁴ + reasonable housing costs	4 for each working family member	75-100%	IT: no SSC: no	45 (if not receiving social assistance)	40%	main benefit: 5-9 per child; child-raising allowance for very young children: additional 5-7	main benefit: none (universal payment); child-raising allowance: 20-40% once income > 62	-	-	-
Greece	-	-	-	-	none at the national level		0.5-1 per child plus additions for large families	reduced in steps for incomes > 65	-	-	-
Ireland	56	19 (only for 2nd earner)	100%	IT: no SSC: no	53	100%	3-4 per child	none (universal payment)	60% of family gross earnings exceeding 88 (higher limit for larger families)	couple jointly working at least 20 hours per week	60% (of gross family earnings)
Italy	-	-	-	-	none at the national level		see employment-conditional benefits		3-17 per family member (also spouses) depending on family type	must work at least 3 days per week; reduced benefits if working less than full-time	reduced in steps for family incomes > 73
Luxembourg	64	13	100%	IT: yes SSC: reduced	6 (must receive social assistance)	100%	8-13 per child; plus education allowance for children aged 3-	none (universal payment)	-	-	-
Netherlands	49	-	100%	yes, but max. amount is shown on net basis	6	54%	2-7 per child	none (universal payment)	-	-	-
Portugal	59	-	100%	IT: no SSC: no	none at the national level		4 per child	reduced to 3 per child once income > 71	-	-	-
Spain	none at the national level				none at the national level		2 for first child, 0.2 for further children	100% once income > 55	-	-	-
Sweden	35 +reasonable housing cost	-	100%	IT: no SSC: no	17	33% (disregard of 18)	4-8 per child	none (universal payment)	-	-	-
United Kingdom	51	2-4	100%	IT: no SSC: no	100% of recognised rent; 100% of council tax	65% (housing benefit); 20% (council tax benefit)	3-5 per child	none (universal payment)	18 + up to 13 per child + 4 if working > 30 hours per week; only entitled if >= 1 child	at least one person working >= 16 hours per week	70% of family income > 29

Abbreviations: IT = income tax; SSC = social security contributions.

Explanatory Notes: Reference year is 1998. Rules for social benefits can vary between regions or municipalities. Where social assistance is subject to job-search or other conditions (e.g. in Denmark), it is assumed that both spouses comply with the relevant requirements. All information is for families with two children. IT = income tax; SSC = social security contributions. Further information, and data for later years, can be found in the EUROMOD country reports (www.iser.essex.ac.uk/msu/emod/documentation/countries/), the OECD Tax Database (www.oecd.org/ctp/taxdatabase) and the OECD series *Benefits and Wages* (www.oecd.org/els/social/workincentives).

¹ in % of median gross earnings of primary earners (not including employer social security contributions)

² cash assistance for privately rented accommodation. Housing benefits may be paid through the social assistance program. In this case, they are already reflected in the social assistance amounts shown in this table.

³ in addition to family-related tax concessions shown in the companion table under income tax. Does not include any benefits available for pregnancy, childbirth, parental leave, or childcare benefits.

⁴ West Germany

References

- [1] Alesina, A. and A. Ichino (2007). “Gender based taxation.” Harvard University Working Paper, April 2007.
- [2] Alm, J. and L. Whittington (1995a). “Income Taxes and the Marriage Decision.” *Applied Economics* 27, 25-31.
- [3] Alm, J. and L. Whittington (1995b), “Does the Income Tax Affect Marital Decisions?” *National Tax Journal* 48, 565-72.
- [4] Alm, J. and L.A. Whittington (1996). “The Rise and Fall and Rise...of the Marriage Tax.” *National Tax Journal*, 49, 571-89.
- [5] Alm, J. and L.A. Whittington (1997). “Til Death or Taxes Do Us Part: The Effect of Income Taxation on Divorce.” *Journal of Human Resources* 32, 388-412.
- [6] Alm, J. and L.A. Whittington (1999). “For Love or Money? The Impact of Income Taxes on Marriage.” *Economica*, 66, 571-89.
- [7] Alm, J., S. Dickert-Conlin and L.A. Whittington (1999). “Policy Watch: The Marriage Penalty.” *Journal of Economic Perspectives*, 13, 193-204.
- [8] Apps, P. and R. Rees (2007). “The Taxation of Couples.” IZA Discussion Paper No. 2910.

- [9] Armstrong, M. and J.-C. Rochet (1999). “Multi-dimensional screening: A user’s guide.” *European Economic Review*, 43, 959-79.
- [10] Bitler, M., J. Gelbach, H.W. Hoynes, and M. Zavodny (2004). “The Impact of Welfare Reform on Marriage and Divorce.” *Demography* 41, 213-236.
- [11] Blundell, R.W. (1995). “The Impact of Taxation on Labor Force Participation and Labor Supply,” in *OECD Job Study: Taxation, Employment, and Unemployment*, Chapter 3, OECD: Paris.
- [12] Blundell, R.W. and T. MaCurdy (1999). “Labor Supply: A Review of Alternative Approaches.” In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Vol. 3A, Amsterdam: Elsevier Science B.V., 1559-1695.
- [13] Boskin, M. J. and E. Sheshinski (1983). “Optimal tax treatment of the family: Married couples.” *Journal of Public Economics*, 20, 281-97.
- [14] Brett, C. (2006). “Optimal Nonlinear Taxes for Families.” *International Tax and Public Finance* 14, 225-261.
- [15] Browning, E.K. and W.R. Johnson (1984). “The Trade-Off between Equality and Efficiency.” *Journal of Political Economy* 92, 175-203
- [16] Bull, N., J. Holtzblatt and J.R. Nunns (1999). “Defining and Measuring Marriage Penalties and Bonuses.” OTA Paper 82.

- [17] Cremer, Helmuth, Jean-Marie Lozachmeur, and Pierre Pestieau (2007). "Income Taxation of Couples and the Tax Unit Choice." Working Paper, 2007.
- [18] Dickert-Conlin, S. (1999). "Taxes and Transfers: Their Effect on the Decision to End a Marriage." *Journal of Public Economics* 73, 217-40.
- [19] Dickert-Conlin, S. and S. Houser (1998). "Taxes and Transfers: A New Look at the Marriage Penalty." *National Tax Journal* 51, 175-217.
- [20] Eissa, N. (1995). "Taxation and Labor Supply of Married Women: The Tax Reform Act of 1986 as a Natural Experiment." NBER Working Paper No. 5023.
- [21] Eissa, N. and H.W. Hoynes (2000a). "Explaining the Fall and Rise in the Tax Cost of Marriage: The Effect of Tax Laws and Demographic Trends, 1984-1997." *National Tax Journal* 53, 683-711.
- [22] Eissa, N. and H.W. Hoynes (2000b). "Tax and Transfer Policy, and Family Formation: Marriage and Cohabitation." Working Paper, December 2000.
- [23] Eissa, N. and H.W. Hoynes (2004). "Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit." *Journal of Public Economics* 88, 1931-1958.
- [24] Eissa, N., H.J. Kleven, and Kreiner, C.T. "Evaluation of Four Tax Reforms in the United States: Labor Supply and Welfare Effects for Single Mothers" *Journal of Public Economics* 92, 2008, 795-816.

- [25] Feenberg, D. and H.S. Rosen (1995). "Recent Developments in the Marriage Tax." *National Tax Journal* 48, 91-101.
- [26] Heckman, J. (1993). "What Has Been Learned About Labor Supply in the Past Twenty Years?" , *American Economic Review Papers and Proceedings*, 83, 116-21.
- [27] Hoynes, H.W. (1997a). "Does Welfare Play Any Role in Female Headship Decisions?" *Journal of Public Economics* 65, 89-117.
- [28] Hoynes, H.W. (1997b). "Work, Welfare, and Family Structure: What Have We Learned?" , in *Fiscal Policy: Lessons From Economic Research*, Alan Auerbach (ed.), MIT Press: Cambridge, Massachusetts.
- [29] Immervoll, H. (2004). "Average and marginal effective tax rates facing workers in the EU. A micro-level analysis of levels, distributions and driving factors." OECD Social, Employment and Migration Working Paper No. 19, OECD: Paris.
- [30] Immervoll, H. and C. O'Donoghue (2003). "Employment transitions in 13 European countries. Levels, distributions and determining Factors of net replacement rates." CESifo Working Paper No. 1091, CESifo: Munich.
- [31] Immervoll H., H.J Kleven, C.T. Kreiner and E. Saez (2007). "Welfare Reform in European Countries: A Microsimulation Analysis." *Economic Journal*, 117, 1-44.
- [32] Kleven, H.J and C.T. Kreiner (2006). "The Efficient Taxation of Couples." Mimeo.

- [33] Kleven, H.J., C.T. Kreiner and E. Saez (2007). “The Optimal Income Taxation of Couples as a Multi-Dimensional Screening Problem.” CESifo Working Paper 2092, September 2007.
- [34] Kleven, H.J., C.T. Kreiner and E. Saez (2008). “The Optimal Income Taxation of Couples.” Working Paper, August 2008. Forthcoming in *Econometrica*.
- [35] Mirrlees, J.A. (1976). “Optimal tax theory: a synthesis.” *Journal of Public Economics* 6, 327-358.
- [36] Mirrlees, James A. (1986). “The Theory of Optimal Taxation,” in K.J. Arrow and M.D. Intrilligator (eds.), *Handbook of Mathematical Economics* vol. 3. Elsevier Science B.V.: Amsterdam.
- [37] Moffitt, R. (1998). “The Effect of Welfare on Marriage and the Family,” in Robert Moffitt (ed.), *Welfare and Family and Reproductive Behavior*, National Academy Press: Washington D.C.
- [38] OECD (2002). *Revenue Statistics: 1965-2001*. Paris: OECD.
- [39] OECD (2003). *National Accounts, Volume II 1990-2001*. Paris: OECD.
- [40] Pechman, J.A. (1987). *Federal Tax Policy*. Brookings Institution: Washington D.C.
- [41] Pechman, J. and G. Engelhardt (1990). “The Income Tax Treatment of the Family: An International Perspective.” *National Tax Journal* 43, 1-22.

- [42] Rosen, H.S. (1977). "Is it time to abandon joint filing?" *National Tax Journal*, 30, 423-8.
- [43] Rosen, H.S. (1987). "The Marriage Tax Is Down But Not Out." *National Tax Journal* 40, 567-75.
- [44] Rubinfeld, J. (1997). "Affirmative Action." *Yale Law Journal* 107, 427-472.
- [45] Saez E. (2002). "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics*, 117, 1039-73.
- [46] Schroyen, Fred (2003). "Redistributive taxation and the household: the case of individual filings." *Journal of Public Economics* 87, 2527-2547.
- [47] Sutherland H. (2001). "Final Report EUROMOD: An Integrated European Benefit-Tax Model." EUROMOD Working Paper No. EM9/01.