NBER WORKING PAPER SERIES

THE GLOBAL LIFE-CYCLE OPTIMIZER – ANALYZING FISCAL POLICY'S POTENTIAL TO DRAMATICALLY DISTORT LABOR SUPPLY AND SAVING

Johannes Brumm
Laurence J. Kotlikoff
Christopher Krause

Working Paper 32335
http://www.nber.org/papers/w32335

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024, revised July 2024

We thank Robert Moffitt, Yongyang Cai, Philip Jung, and seminar participants at Boston University, the European Central Bank (ECB), the Federal Reserve Bank of Atlanta, and at Karlsruhe Institute of Technology for helpful suggestions. Johannes Brumm and Christopher Krause acknowledge support from the ERC (101042908). Laurence Kotlikoff acknowledges support from the Federal Reserve Bank of Atlanta, Economic Security Planning Inc., the Goodman Institute, and Boston University. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Global Life-Cycle Optimizer – Analyzing Fiscal Policy's Potential to Dramatically Distort Labor Supply and Saving
Johannes Brumm, Laurence J. Kotlikoff, and Christopher Krause
NBER Working Paper No. 32335
April 2024, revised July 2024
JEL No. H2,H3,H30,H31,I38,J22

ABSTRACT

Fiscal policy in the U.S. and other countries renders intertemporal budgets non-differentiable, non-convex, and discontinuous. Consequently, assessing work and saving responses to policy requires global optimization. This paper develops the Global Life-Cycle Optimizer (GLO), a stochastic pattern-search algorithm. The GLO robustly, precisely, and quickly locates global optima in highly complex fiscal settings. We use the GLO to study how a stylized U.S. fiscal system distorts workers' labor supply and saving assuming standard preferences. The system incorporates kinks from federal personal income tax brackets, Social Security's FICA tax, and a notch from the provision of basic income below a threshold. The GLO reproduces theoretically predicted earnings bunching and flipping over a remarkably wide range of wage rates. Saving distortions can be equally dramatic. Associated excess burdens range from substantial to massive. Restricting labor supply to full- or part-time work can eliminate flipping when it's optimal and produce flipping when it's sub-optimal. Joint filing can significantly reduce the earnings of lower-wage spouses relative to that of higher-wage spouses. The GLO can be applied to assess a country's or state's full set of work and saving disincentives. Consequently, it can facilitate analyses of structural labor supply and tax reform.

Johannes Brumm
Karlsruhe Institute of Technology
Waldhornstr. 27
76131 Karlsruhe
Germany
jo.brumm@gmail.com

Laurence J. Kotlikoff
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
kotlikoff@gmail.com

Christopher Krause
Waldhornstr. 27
Karlsruhe 76131
Germany
christopher.krause@kit.edu

1 Introduction

Fiscal systems in many countries are extraordinarily complex thanks to a plethora of national and state tax and benefit policies. These policies generally come with program-specific eligibility conditions and highly non-linear net payment schedules. Both sets of provisions can depend on a range of economic and demographic factors. These include labor income, asset income, total income, wealth, purchases of particular goods and services, marital status, and the presence of children. Kinks in budget sets from changes in tax brackets are routine. So are notches arising from benefit eligibility and tax thresholds.[1] Furthermore, choice sets are intertemporally intertwined, as future net tax schedules are endogenous to current decisions. Thus, saving more this year means more future asset income. This can place a worker in higher future income-tax brackets or leave them ineligible for future income-tested benefits. The prevalence of highly complex, intertemporal budget constraints is no surprise. Politicians signal their value added by enacting new policies. In so doing, they rarely consider the impact on work or saving incentives let alone economy-wide efficiency. The U.S. is a case in point. Its fiscal system features over 1000 tax and benefit programs comprising federal programs, state programs, and 51 state-specific versions of federal programs.

Understanding the effects of fiscal complexity falls on economists. The challenge is handling non-differentiable, non-convex, and discontinuous (NND) intertemporal budget sets. In such environments, work and saving responses cannot be assessed by relying on first-order optimality conditions. Instead, they require global optimization to identify the best of all affordable age-consumption and age-leisure paths. But feasible global-solution methods must avoid the *curse of dimensionality*, i.e. computation requirements rising exponentially in the dimensionality of the problem – in this case, the large number of continuous choices made over the life cycle.[2] This paper develops the *Global Life-Cycle Optimizer* or GLO. The GLO searches for globally optimal annual consumption and labor-supply paths building on the pattern search literature (see Torczon, 1997; Audet and Dennis, 2003). In each iteration, the GLO considers several potential improvements to the current guess. These are generated by adjusting a randomly chosen subset of consumption and labor supply decisions while imposing lifetime budget balance and penalizing violations of cash-flow constraints. The algorithm considers ever larger steps when improvements are found, ever smaller steps when none are found, and stops when the step size falls below a specified minimum.

Given the GLO's simplicity, its performance is remarkable. We first introduce the GLO and illustrate its speed, accuracy, and robustness. Next we consider a stylized NND net-tax schedule comprising just three elements – the U.S. federal income-tax brackets, Social Security's FICA tax, and basic income support. Income support is limited to those earning less than a threshold amount set at roughly half the annual earnings of minimum-wage workers.[3] Depending on the worker's wage, this modest set of fiscal provisions can produce earnings *bunching* and earnings *flipping* over remarkably wide ranges of the earnings distribution. Earnings bunching entails earning just below higher marginal tax and benefit-eligibility thresholds. Earnings flipping references supplying small amounts of labor in some years and high amounts in others. These policy-induced major labor supply reductions dramatically lower annual saving and, consequently, wealth at retirement. Earnings bunching has long been theoretically predicted (see Moffitt, 1983; Burtless, 1976) and empirically documented (see Kotlikoff, 1978; Friedberg, 1998, 2000). More recently, economists have used bunching behavior to

[1] E.g., U.S. Medicaid, Supplemental Security Income programs, Section-8 housing thresholds, and Medicare Part B IRMAA premium thresholds.

[2] A different form of the curse of dimensionality is often encountered in dynamic stochastic economic models, even absent NND constraints, when the number of continuous *states* becomes large (see, e.g., Brumm and Scheidegger, 2017).

[3] This notch proxies for the complete loss of Medicaid, Section-8 housing, Supplemental Security Income, and other benefits from earning beyond specified limits.

infer labor supply elasticities (see, e.g., Saez, 2010; Chetty et al., 2011a), with other economists questioning whether these measures are properly identified (see Blomquist et al., 2021). Earnings flipping across years is also the theoretically expected response to NND frontiers. It permits workers to partially convexify their lifetime budget sets – to work less and pay lower taxes, on average, over one's workspan. Here, again, evidence supports theory. Gustman and Steinmeier (1983, 1984) report that a substantial share of workers spend major portions of their working lives in part-time work. These and others studies indicate substantial heterogeneity in "retiring," with some workers retiring gradually, others doing so abruptly, and many retiring and then un-retiring, typically to part-time work (see Rust, 1989).

To demonstrate GLO's versatility, we consider discrete labor choice, joint taxation of married couples' earnings, and taxation of total income. Discrete choice references limiting the choice set to working either what we define as full time or what we define as part time in a given year. Doing so can lead to flipping when it would otherwise not arise. It can also eliminate flipping when it would arise were the worker unconstrained. Hence, restricting workers to part- or full-time work is hardly a benign assumption. Taxing married couples jointly has the theoretically predicted (see, e.g., Kaygusuz, 2010; Guner et al., 2012; Bick and Fuchs-Schündeln, 2018) impact of discouraging the relative labor supply of lower-wage spouses. Here, as elsewhere, the GLO not only evinces what theory predicts and data confirm. It also provides a quantitative sense of fiscal distortions. Our third extension is taxing total income, i.e., asset as well as labor income. This intertemporally intertwines taxes in the present with those in the future. It not only, as expected, alters the age-consumption profile. It can also limit current earnings and saving to limit future taxes.

1.1 Our Stylized Fiscal System and Assumptions on Preferences

Our net-tax schedule comprises three elements. The first is a wage tax that includes the seven brackets of the U.S. 2022 federal personal income-tax. Table 2 reports these brackets and their associated marginal tax rates. Figure 2 depicts these brackets by plotting after-tax against pre-tax income.[4] The second is the 12.4 percent Social Security payroll (FICA) tax levied up to its 2022 \$147,000 ceiling ignoring, as is standard (see Burtless, 1976), marginal Social Security benefit-tax linkage.[5] The third element is the provision of \$10,000 in basic income for those earning less than \$15,000. Figure 1 shows a static budget constraint (ignoring saving), including all three elements.

Our time-separable preference structure is drawn from the standard King et al. (1988) class. Period-specific utility is the log of consumption less a labor-supply parameter times an isoelastic function of labor supply. Agents begin work at 25, retire at 65, and die at 85. Given our chosen equal interest and time preference rates, optimal age-consumption profiles are flat absent binding cash-flow constraints. This holds despite the nature of wage taxation. The optimal age-labor supply profile is also flat when real wages are constant and workers face either lump-sum taxes or linear wage taxes. Consequently, in most of our exercises, the shapes and levels of workers' age labor-supply profiles convey, at a glance, policy distortions. Absent non-linear taxes, our preferences dictate identical constant labor supply through retirement regardless of the worker's wage rate, i.e., income and substitution effects fully offset.

[4] Our stylized policy taxes, for much of the paper, only labor income. This isolates fiscal impacts on labor supply and, given our posited preferences, associated changes to saving. Focusing on labor-income taxation is for illustrative purposes. The GLO is fully capable of handling total-income taxation, as we show in section 6.

[5] This assumption is reasonable given Social Security's immense complexity. The system provides just 12 benefits, but has 2728 basic rules about the 12 benefits in its Handbook. And its Program Operating Manual System has 20,000 pages of rules about the 2728 rules about the 12 benefits.

Figure 1: Stylized Static Budget Constraint with Kinks and Notch

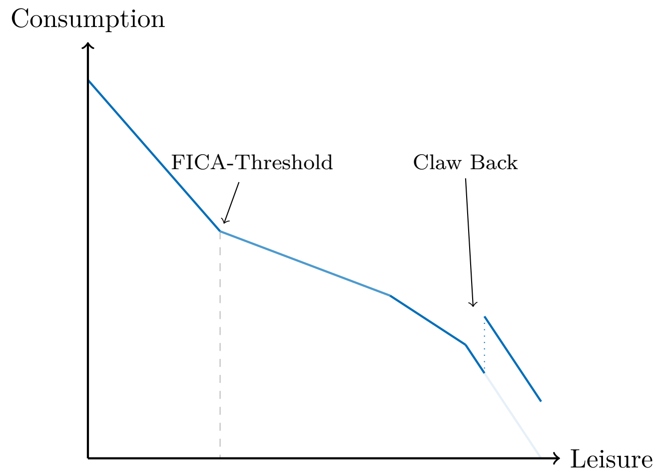

Note: This static budget constraint includes a convex kink from the FICA tax, two concave kinks from changes in income-tax brackets, and a notch from the claw-back threshold for basic income. For illustrative purposes, this figure abstracts from the intertemporal dimension of the life-cycle optimization problem.

1.2 Overview of Findings

Our principle finding is simply the ability of the GLO to determine global solutions quickly, precisely, and robustly in complex settings. Quickly refers to running young households with 60-year lifespans to completion in minutes. Precisely references recovering analytically determined solutions to, in most cases, five decimal places. Robustly references finding the same optimum regardless of the GLO's starting guess.

As indicated, the GLO reproduces theoretically-predicted and empirically-documented labor supply behaviors. What's surprising is the wide range of wages over which it occurs. One would expect major bunching around the $15,000 earnings level to avoid loss of basic income. But it also occurs for workers with quite high levels of wages who work less to lower their tax bracket from the 32 percent to 24 percent. The impact on labor supply from earnings bunching can be massive – up to 51.1% in our setting. This change is relative to labor supply under equal annual lump-sum taxation. For workers whose wages are fixed through time, earnings bunching, when it arises, generally entails setting their labor to a fixed lower level. For workers with growing real wages, earnings bunching entails working less each year. This occurs until wages grow sufficiently to reach a threshold. At this point, it pays to work more and, thereby, take advantage of the higher gross wage despite moving into a higher tax bracket.

In addition to permanently lowering their labor supply, workers, depending on their wage rates, will flip back and forth between high and low labor supply. Flippers may switch between earning just before the notch or a kink in some years and working substantially more in other years, notwithstanding facing higher marginal net taxation. Workers near the FICA kink are particularly prone to flip. Since the FICA kink is concave, it induces the opposite of bunching namely a desire to spend part of one's working years working intensively above the Social Security taxable earnings ceiling. The reason is to garner a higher net wage. Obviously, this behavior depends on preferences. A higher or lower time-preference rate would, for example, change the range of wage rates that admit particular behaviors. More interestingly, as we show, small differences in time preference for a given interest rate can produce dramatically different timings of labor supply for workers who flip. Stated differently, the precise timing of labor supply can be hypersensitive to parameter values. Moreover, for a given set of parameter values, the timing of flipping appears quite arbitrary and this timing is highly sensitive to the worker's wage rate. Another key finding is the extreme sensitivity of life-cycle saving to the fiscal system. Wealth at retirement is, in our worst-case finding, 40 percent lower than it would be with equal-annual lump sum taxes.

Flipping between part-time and over-time work may also be optimal for workers with only those two options. But, as we show, if one imposes discrete labor-supply choice for computational convenience, i.e., when such constraints do not, in fact, exist, the resultant constrained optimum can deviate dramatically from the unconstrained optimum. For example, constraining agents to work either part- or full-time when they aren't so constrained can lead to flipping when it would otherwise not arise. It can also preclude flipping when flipping would arise. This is the first of three extensions to our primary analysis. The second studies labor supply of married couples subject to joint taxation. As theory predicts, joint progressive taxation can induce more work by the higher wage spouses and less work by the lower wage spouse. Our third extension considers total income rather than only labor income taxation. As we show, taxing total income distorts the age-consumption profile and can lead to flipping for part of one's working life, but not the rest.

We also find moderate to exceptionally large excess burdens measured relative to lump-sum taxation. These values are highly sensitive to workers' wages and fiscal provisions. With all three fiscal elements in place, low-wage workers experience excess burdens ranging as high as 26 percent. This reflects the presence of the notch at \$15,000 leading, for example, a worker with a full-time wage of \$30,000 to earn, annually, only \$15,000 to forego sacrificing the otherwise available \$10,000 in basic income. Absent the basic-income-cum-clawback policy, excess burdens are far smaller. In the case of the \$30,000 full-time wage worker, excess burden with all fiscal elements is 26 percent. Absent the provision of basic income with its clawback notch, excess burden is only 2 percent. Given the full policy's massive distortion, does it leave the \$30,000 worker significantly better off? The answer, surprisingly, is yes. The worker's welfare ends up 27 percent higher. While the worker's annual consumption falls by 19 percent their annual labor supply falls by 55 percent. Still, inducing the poor not to work may have major social costs not included in this model.[6]

Table 1 conveys our most important bottom line. It shows that even a relatively simple tax system will lead workers, across a wide range of wage rates, to modify their labor supply by bunching or flipping. These are, of course, extreme behaviors. They arise when consumption and leisure are reasonably good substitutes – our base case. But they also occur when our assumed Frisch elasticity is relatively low. This would be surprising were the fiscal system simply rotating linear budget frontiers. But they are dramatically altering the shapes of budget frontiers, inducing discrete rather than marginal changes in behavior.

1.3 Related Literature

The early literature considering responses to kinks and notches (aka cliffs) includes Kotlikoff (1978), Zabalza et al. (1980), Danzinger and Plotnick (1981), Moffitt (1983), Hausman (1985), Pencavel (1986), Fraker and Moffitt (1988), Rust (1989), Moffitt (1992), Hoynes (1996), Hagstrom (1996), Keane and Moffitt (1998), Blundell and MaCurdy (1999), Eklöf and Sacklén (2000), Meyer and Rosenbaum (2001), Moffitt (2002, 2003), and Blundell and Hoynes (2004). More recent studies, including Saez (2010), Chetty et al. (2011a), Brown (2013), Bastani and Selin (2014), Blomquist et al. (2021), and Bertanha et al. (2023), further document bunching, use observed bunching to

[6]This potentially includes not getting married due to the loss of benefits that doing so would entail. See, in this regard, Ilin et al. (2022). This, in turn, means a higher share of children growing up, in poverty with a single parent.

Table 1: Labor Supply Behavior

Wage Range	Features of Labor Supply
$- \$30k$	bunching at claw back threshold
$\$30k - \$45k$	flipping between claw back threshold and above
$\$45k - \$82k$	no flipping or bunching
$\$82k - \$83k$	bunching at convex kink at $89,075
$\$83k - \$123k$	no flipping or bunching
$\$123k - \$135k$	flipping around concave FICA kink
$\$135k - \$144k$	no flipping or bunching
$\$144k - \$152k$	bunching at convex kink at $170,050
$\$152k - \$193k$	no flipping or bunching
$\$193k - \$196k$	bunching at convex kink at $215,950
$\$196k -$	no flipping or bunching (except at $539k kink)

estimate labor supply elasticities, and question whether bunching data suffices to identify underlying behavioral parameters.

Of most relevance to our paper are those recognizing the need for global search and implementing global search routines. Burtless (1976)'s seminal paper suggested using piece-wise linear budget constraints (PLB) to compare utility along different segments of non-convex budget frontiers. Comparing, for given preference parameters, indirect utility along each constraint as well as direct utility at kink points revealed the global optimum. Unfortunately, the PLB approach becomes computationally infeasible when expanded to multiple fiscal programs let alone multiple periods. Consequently, Burtless (1976), Friedberg (1998), Friedberg (2000), and others using PLB do so primarily within static (one-period) models featuring a limited number of fiscal programs. PLB's computational constraints as well as the presence, in the data, of households making theoretically dominated labor supply decisions led MaCurdy et al. (1990) to approximate kinked and notched frontiers with smoothed functions. In so doing, they were able to incorporate wage-rate measurement error, which could explain observed dominated choices as well as permit testing rationality (Slutsky) conditions.

Yet, smooth frontiers can't explain earnings bunching or discrete changes in labor supply from year to year. Zabalza et al. (1980), Fraker and Moffitt (1988), Keane and Moffitt (1998), take a different approach to computing global optima of structural models. They drastically restrict the choice set. Specifically, households must consume their current income, i.e., they can neither save nor dis-save. And labor supply is restricted to zero work, part-time work, or full-time work in static contexts. Blundell and Shephard (2012) examine UK tax reform, including the use of tagging and imperfect observations on hours work. Theirs is also a static model, but incorporates six different discrete choices of labor supply.

Assuming single-period agents facing discrete labor-supply constraints comes at a price. Clearly, most households do save or dis-save. Yes, many households, particularly those with low incomes, are cash-flow constrained. But the degree to which their constraints bind is endogenous to future household choices. And restricting labor supply to three discrete options can rule out or rule in suboptimal behavior when such constraints don't apply. Indeed, we offer an example in which discretizing labor supply can, depending on the definition of part-time and full-time, lead to flipping when it otherwise would not arise. It can also preclude flipping when it would otherwise arise. Defining full-time and part-time work also begs the question of measurement. Full-time work for newly minted lawyers and doctors is generally 60 or more hours per week, not 40 hours. As for part-time work, there are many means of working fewer hours during a year.[7] Expanding the discrete choice set limits concern about mis-specifying the options. Blundell and Shephard (2012) report that specifying more than six options makes little difference to their results. The literature, e.g., Keane and Moffitt (1998), also stresses the need to incorporate welfare stigma given the high degree of non-participation in particular public-assistance programs. Our illustration of GLO abstracts from this critical concern, but including stigma in preferences seems eminently feasible.

Like Keane and Moffitt (1998), Rust (1989, 1990) pursues global analysis via discretization. But Rust proposes incorporating intertemporal choice by implementing discrete-state dynamic programming. He also limits consumption and labor supply to discrete values. However, he too readily acknowledges computational limits. For example, fully considering Social Security requires treating all past covered earnings as state variables. The GLO avoids this concern that renders high-dimensional dynamic programming infeasible due to the curse of dimensionality. Rust sidesteps this problem by assuming Social Security benefits are determined by average past-covered earnings. These studies incorporate uncertainty to facilitate realistic empirical analysis. Our goal is far simpler – clarifying, based on assumed behavioral parameters, the nature of potential responses to NND policy. Stated differently, we seek to describe/characterize complex fiscal policies in terms of the responses they can induce assuming benchmark preferences. This is a very different objective from measuring the precise responses of workers to specific policies.

Moore and Pecoraro (2020, 2021, 2023) also restrict labor supply to discrete values in doing global life-cycle optimization. But they provide four major improvements over prior discrete-choice analyses. First, they incorporate an extensive range of NND fiscal policies. Second, their analysis is general equilibrium, entailing computation of the transition path of life-cycle economies featuring NND policies. Third, they incorporate a wide range of realistic elements, including housing choice, the decision to rent or own, fixed costs of working, home production, child care costs, bequests, uncertain longevity, and more. And fourth, building on Carroll (2006) and Iskhakov et al. (2017), they introduce a hybrid endogenous grid method approach to dynamic programming that avoids interpolation and permits global optimization. Their novel approach specifies a reduced household state space so that exact interior solutions to all of the time-t continuous choice variables can be computed over a single exogenous grid of time $t+1$ net worth, for each possible combination of the time t discrete labor and residential status choice variables. The global solutions are identified from a set of local optima in each period by application of the discrete-choice endogenous grid method of Iskhakov et al. (2017).[8] Consumption choices that violate cash-flow and other rigid constraints are ruled out by penalizing candidate value functions. Moore and Pecoraro's papers teach valuable lessons. The most important for our analysis is their finding that incorporating explicit NND fiscal policy materially matters both to microeconomic behavior and macroeconomic outcomes.

Before developing the GLO as a robust global optimizer for deterministic NND life-cycle models with continuous choice, we tried a variety of traditional global optimization methods, including genetic and random search algorithms – without success. Building on the pattern search literature (see Torczon, 1997; Audet and Dennis, 2003) and including features tailored to fiscally realistic life-cycle problems proved successful. Of course, other methods, including those discussed and developed in Arnoud et al. (2019) and Guvenen (2011), may be able to match GLO's optimization performance if properly adapted to handle such fiscal conditions. Arnoud et al. (2019) provide a comparison of specific global optimizers, including several versions of TikTak developed in Guvenen (2011), which

[7] The list includes working less intensely per hour, working fewer hours per day, working fewer days per week, and working fewer weeks per year. Hence, setting part-time work to the same, say, 20 hours per week for 52 weeks for all workers may be problematic.

[8] This method constructs the upper envelope of time $t+1$ value functions. The maximand of time-t lifetime utility across all time $t+1$ potential grid points provides the value function at each time-t grid point.

they apply to method-of-simulated-moments estimation as well as to analytical test functions. We also test the GLO against these functions, highlighting its speed, reliability, and scalability.

One might also ask whether neural nets could help find global solutions in our setting. Azinovic et al. (2022), Maliar et al. (2021), and Duarte et al. (2021) demonstrate the impressive power of machine learning to handle complex life-cycle problems, including various forms of uncertainty. Yet, as discussed in Duarte et al. (2021), machine learning is ill suited for dealing with NND frontiers due to its reliance on differentiability.

1.4 Organization of this Paper

The paper is organized as follows. Section 2 lays out our life-cycle optimization problem and details our tax system. Section 3 presents GLO and demonstrates its accuracy. Section 4 scrutinizes the response of labor supply to kinks and notches in the tax code. Section 5 reports the quantitative impact of different tax schemes on labor supply and wealth at retirement; it also measures and decomposes our tax system's excess burden. Section 6 presents extensions, and section 7 concludes.

2 Our Life-Cycle Problem

Our life-cycle model is simple apart from assumed kinks and notches in the net tax schedule. Depending on the case under consideration, income will reference either wage income or total (wage plus asset) income.

2.1 Lifetime Utility

Households live for T periods with lifetime utility given by

$$\sum_{t=1}^{T}\left(\frac{1}{1+\rho_t}\right)^{t-1}U(c_t,l_t), \tag{1}$$

where ρ_t is the (potentially time-varying) time preference rate, c is consumption, and l is labor supply. We assume that per-period (annual) utility obeys the commonly-applied King et al. (1988) additively separable functional form,

$$U(c,l)=\log c-\chi\frac{l^{1+1/\gamma}}{1+1/\gamma}, \tag{2}$$

where γ is the Frisch elasticity and χ is a scaling parameter, which is set to 1 for most of the analysis.

Households work for the first R periods only, thus $l_t = 0$ for $t > R$. The per-period budget constraint in t is:

$$c_t + a_{t+1} = w_t l_t + (1+r_t)a_t - \mathcal{T}(y_t), \tag{3}$$

where a_t are the beginning-of-period assets, w_t is the wage rate, r_t is the interest rate on assets, $\mathcal{T}_t$ is the net-tax function, and $y_t = w_t l_t$ is labor income.[9] Note that if the after-tax wage rate is constant and the interest rate equals the time preference rate, $r_t = \rho_t$, labor supply is constant and

[9]We consider taxing total income, i.e. $y_t = w_t l_t + r_t a_t$, in section 5.

independent of γ. We focus on examples satisfying this assumption since it implies that any deviation from a constant age labor-supply profile reflects the kinks or notches of our net tax schedule.

Throughout the paper, we set the period length to be one year, and for the most part, we make the following parameter choices. We set the annual interest rate on assets and the time preference rate to 2 percent. For the Frisch elasticity, we choose a value of 1, which is among the higher range of values found in the empirical literature (see Chetty et al., 2011b). We assume that households in the model have an economic lifespan of 60 years, in which they work for 40 years. Finally, agents are not able to borrow, that is $a_{t+1} \geq 0$ for all t.

2.2 Our Stylized Tax System

We consider net-tax codes with n tax brackets, each characterized by cutoff value b, intercept x, and proportional tax rate τ:

$$\mathcal{T}(y_t) = \begin{cases} x_1 + t_1 y_t & \text{if } y_t \in [0, b_1] \\ x_i + t_i(y_t - b_{i-1}) & \text{if } y_t \in]b_{i-1}, b_i], \quad \text{for } i = 2, \ldots, n-1, \\ x_n + t_n(y_t - b_{n-1}) & \text{if } y_t \in]b_n, \infty[. \end{cases}$$

Our baseline net-tax schedule comprises three components. First, the 2022 US single's tax brackets reported in table 2. Second, Social Security's 12.4 percent FICA tax on labor earnings through \$147,000.[10] Third, a basic income of \$10,000 paid to those with incomes below \$15,000 – a proxy for welfare-benefit thresholds arising under the Supplemental Security Income, Medicaid, Section-8 Housing, and other programs. In addition to considering this baseline schedule we also present results assuming just income taxation, $\mathcal{T}_I$, and just income and FICA taxation, $\mathcal{T}_{IF}$. When comparing the three schedules we denote the baseline tax schedule by $\mathcal{T}_{IFB}$ and refer to the three schedules as "INC.", "INC.+FICA", and "INC.+FICA+BASIC". Note that the income-tax scheme has convex kinks only, that the FICA tax adds a non-convex kink at the FICA threshold (see figure 3), and that the baseline tax code exhibits a notch at the claw-back threshold for basic income (see figure 4).

Table 2: Tax Brackets and Marginal Tax Rates

Taxable Income	Marginal Tax Rate
\$0 – \$10,275	10%
\$10,275 – \$41,775	12%
\$41,775 – \$89,075	22%
\$89,075 – \$170,050	24%
\$170,050 – \$215,950	32%
\$215,950 – \$539,000	35%
\$539,000+	37%

Note: Tax brackets for single filers in 2022, see www.forbes.com.

[10] There is no ceiling on the 2.7 percent FICA tax. In what follows, the FICA tax denotes only the Social Security portion of the overall FICA tax. We also treat the division of "employer" and "employee" portions of the FICA tax as economically irrelevant; i.e., we assume workers bear the full tax regardless of how the remittance of the full tax by employers to the government are labeled. The notch can be viewed as proxying for welfare-benefit thresholds arising under the Supplemental Security Income, Medicaid, Section-8 Housing, and other programs.

Figure 2: Income Tax – Convex Kinks Only

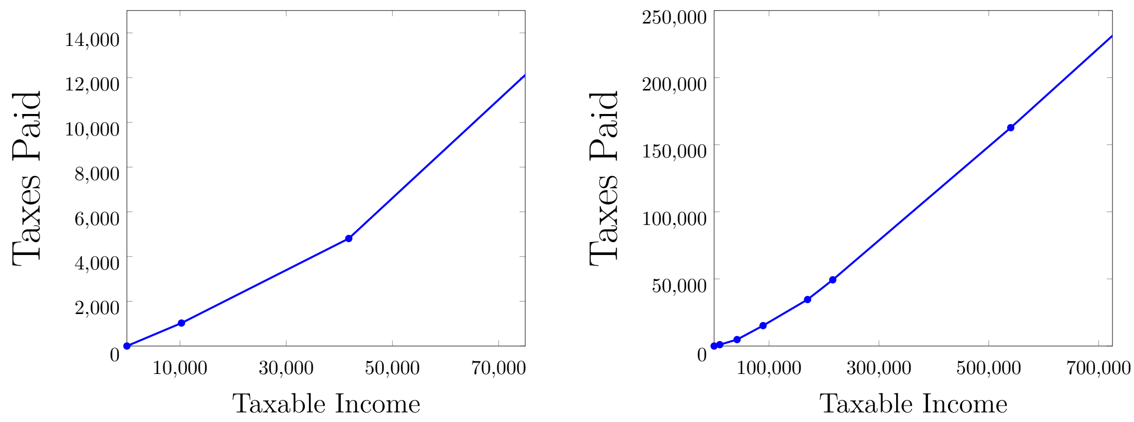

Figure 3: Income Plus FICA Tax – Convex and Concave Kinks

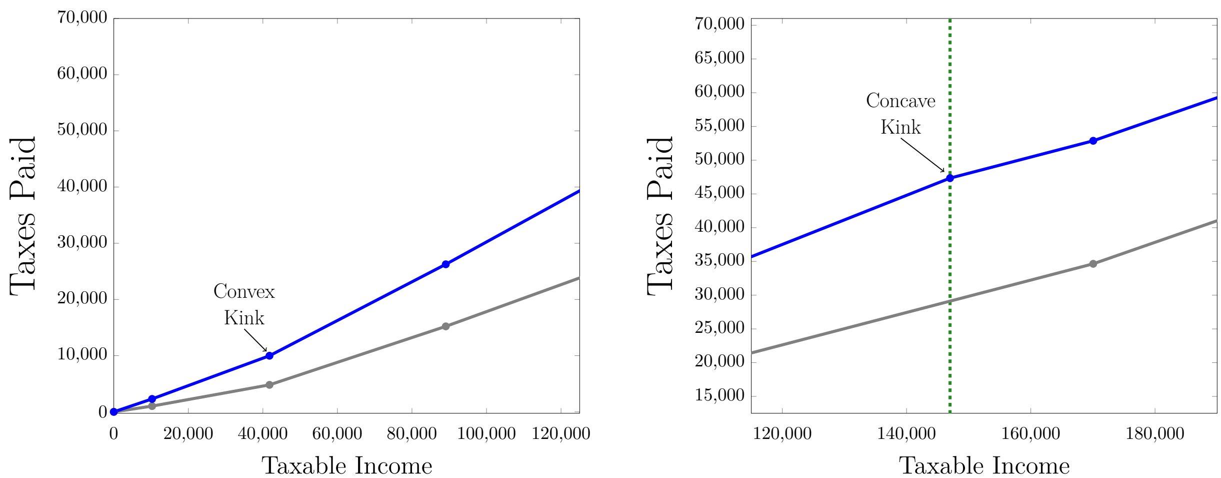

Figure 4: Income Plus FICA Tax Code With Basic Income – Notch at Claw-Back Threshold

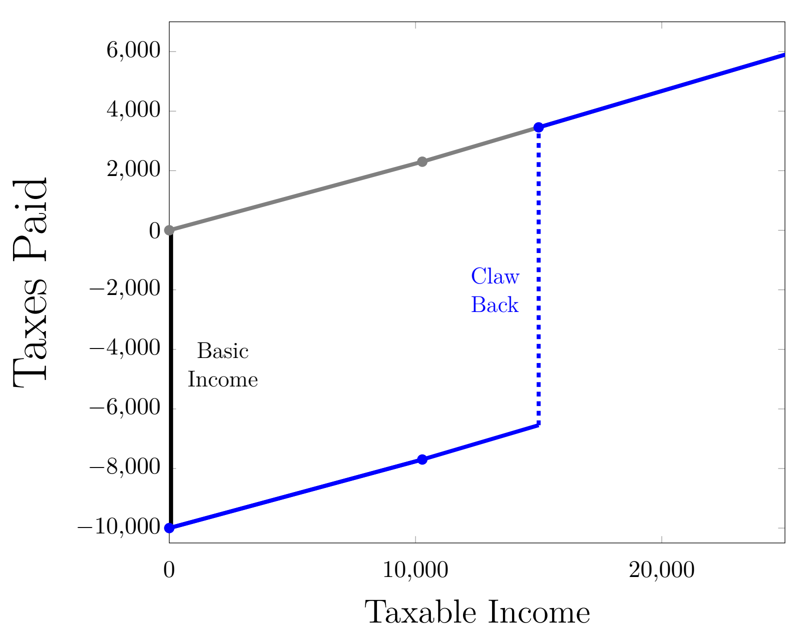

3 The Global Life-Cycle Optimizer

This section presents the GLO algorithm and then showcases its ability to solve highly non-convex life-cycle problems.

3.1 The Algorithm

The GLO algorithm is a variant of pattern search – a well established global optimization method.[11] The basic algorithm seeks the global maximum of a function $f : \mathbb{R}^N \to \mathbb{R}$. Starting from an arbitrary point, $x^0 \in \mathbb{R}^N$, it generates a sequence of points that weakly improve the objective function, $f(x_{k+1}) \geq f(x_k)$. To generate the next point in the sequence, x_{k+1}, the algorithm carries out a so-called polling step.[12] In each polling step a set of J points, the poll set, is constructed by moving from the current point in J different directions, i.e. adding J different vectors to x_k. Each of these J vectors is given by multiplying a spanning direction $d \in D = \{d^1, \dots, d^J\}$ with a scalar m, called the mesh size.[13] The poll set is thus given by:

$$\mathcal{P} = \{x_k + m \cdot d : d \in D\} \tag{4}$$

If one of the points in the poll set improves the objective function relative to the current point, $f(p) > f(x_k)$, the poll is called successful and the point that achieves this improvement becomes the current point in the next iteration, $x_{k+1} = p$.[14] If no improvement is found, $f(p) \leq f(x_k)$ for all $p \in \mathcal{P}$, then the current point is retained for the next iteration, $x_{k+1} = x_k$. Furthermore, if the poll was successful, the mesh size is increased such that the algorithm "zooms out" and considers a larger space to find further improvements. In the case of an unsuccessful poll, the mesh size is reduced, i.e., it "zooms in." The procedure continues until the mesh size falls below a given threshold. Our life-cycle problem consists of finding the agent's paths of consumption and labor supply that maximize lifetime utility subject to per-period cash-flow constraints.

To apply pattern search to this specific optimization problem, we make three substantial adjustments to the basic algorithm. First, when constructing elements of the poll set we simultaneously consider changes in three, not one, dimension. We draw two of these dimensions randomly permitting trade offs between any two values of consumption and/or labor supply. In addition, we always adjust terminal consumption to ensure lifetime budget balance, that is, all remaining assets are consumed in the last period. The second modification entails making the change in one of the chosen dimensions stochastic, replacing the mesh size with a random number drawn from a uniform distribution on an interval around the mesh size. As a third adjustment to the basic algorithm, we include (quadratic) penalty terms for points that violate the cash-flow constraints. Our objective function is thus realized lifetime utility minus the sum of quadratic loss functions. Thanks to this combination of features, GLO's algorithm can, as now shown, tackle highly non-convex life-cycle problems.

We provide a formal description of the GLO algorithm in appendix B. Our application of the GLO to our specific life-cycle NND problem proceeds as follows. Recall, agents work from age 25 through age 65 and consume from age 25 through age 85. Consequently, their optimal solution comprises a

[11] See Torczon (1997) or Audet and Dennis (2003) for a general exposition.

[12] The general pattern search algorithm can also include a so-called search step in addition to the polling step, see Audet and Dennis (2003).

[13] A common pattern for the spanning directions is to vary only one dimension at a time, yet in both directions, thus setting $J = 2N$ with $D = \{(1, 0, \dots, 0), ..., (0, 0, \dots, 1), (-1, 0, \dots, 0), ..., (0, 0, \dots, -1)\}$.

[14] If one chooses a complete polling, the point associated with the best objective function value is chosen. Alternatively, one could stop the polling after any improvement is found and go to the next iteration.

100-element vector with 60 annual levels of consumption and 40 annual levels of labor supply. The GLO starts by setting a mesh size and choosing, at random, 99 of the 100 elements. The 100th element references age-85 consumption, which is always set to satisfy the household's intertemporal budget given values of the other elements. Call this vector X and denote the value of our maximand – lifetime utility plus the sum of 99 cash-flow penalty functions – $f(X)$. The GLO next chooses two of the first 99 elements at random. The first of the two values is both increased and decreased by the mesh size. The second of the two values is both increased and decreased by the current mesh size multiplied by a number randomly chosen from a uniform distribution between 0 and 1.25. The four up-up, up-down, down-up, down-up perturbations together with the other 96 vector elements (including the lifetime-budget balancing final consumption level) provide four candidate solution vectors. We evaluate our objective function for each and repeat the process 500 times. We then find the maximum of the 2000 (4 times 500) evaluations. Call the maximizing vector Y. If $f(Y) > f(X)$, set $X = Y$ and double the mesh size. If $f(Y) < f(X)$, leave X unchanged and reduce the mesh size in half. Next, we repeat the above process starting with the potentially revised X until the mesh size falls below a specified value. Finally, we restart the entire algorithm 200 times with either different starting guesses and/or different seeds for the random number generator, and take the maximum of the stored optima as our solution.

Why does the GLO find the global optimum? Intuitively, if the candidate solution vector is far from the optimum, the algorithm will move to a vector that is close to a local or global optimum. Near a local optimum, large and small deviations (the mesh size) may succeed. Small deviations that succeed result in an increase in the mesh size and thus move the solution away from the local optimum. Near a global optimum, small deviations may succeed, bringing the solution even closer to the global optimum. Large deviations, in turn, will be unsuccessful resulting in an ever smaller mesh size eventually triggering the convergence criterion, which ends the algorithm. The combination of initializing the GLO with a random guess and incorporating multi-starts delivers the global optimum with, as now shown, high accuracy.

3.2 Testing the Algorithm

Figure 5 illustrates the convergence of GLO for the case of preference parameters ensuring, based on first-order conditions, perfectly flat age-consumption profiles. The figure shows, to the far left, our initial guess of these profiles. The subsequent columns display these profiles after 200, 400, 1000, and 5700 iterations of the algorithm, respectively. From left to right, the consumption path becomes smoother and smoother, finally reaching a completely flat profile – the analytic optimum. When it comes to the labor profile, going from left to right, the algorithm increasingly 'discovers' that it is optimal to stay either below the basic-income threshold or to work substantially more, i.e., lifetime utility maximization entails flipping between these two levels of labor supply. The algorithm meets the GLO's convergence criterion after 5723 iterations, running for several minutes.

With fiscal kinks and notches, neither analytic solutions nor convergence proofs are available to test GLO's accuracy. Hence, we assess accuracy by examining, in table 3, the difference in solutions associated with 200 different starting guesses. Specifically, we compute, for all 200 solutions, the percentage deviation in consumption equivalent variation (CEV), as defined in appendix A, from the best solution. Table 3 reports various percentiles of this distribution and does so for four different wage rates. For instance, we find that, for the high-wage case – an annual wage rate of \$130.000, the median solution deviates from the best solution by only 0.0001%, measured as a CEV. This is despite the fact that there is substantial labor-supply flipping due to the concave kink induced by FICA as discussed below. The GLO is less accurate when the wage rate is \$40,000 – a low enough wage to make

Figure 5: Convergence Process – Life-Cycle Profile From Starting Guess to Solution

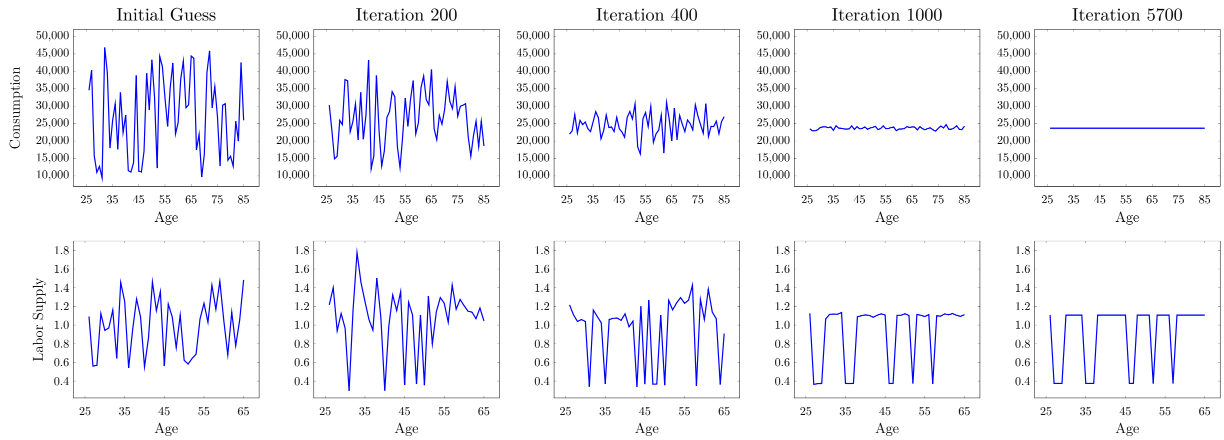

Note: First column shows consumption and labor supply profiles of initial guess. Subsequent columns display these profiles after 200, 400, 1000, and 5700 iterations.

Table 3: Accuracy Assessment – Dependence of Objective From Starting Guess

Wage	1%	10%	Median	90%	Max	Std
20,000	−0.0000	−0.0000	−0.0000	−0.0000	−0.0000	0.0000
40,000	−0.0023	−0.0362	−0.0779	−0.1414	−0.2718	0.0428
80,000	−0.0000	−0.0000	−0.0000	−0.0000	−0.0000	0.0000
130,000	−0.0000	−0.0000	−0.0001	−0.0005	−0.0013	0.0002

Note: Statistics for solving the life-cycle problem with 200 different starting guesses. Percentiles for the percentage deviation in CEV from the best solution are reported. E.g. the median solution of the $40,000 wage-rate problem has a CEV of 0.0779 percent less than the CEV of the best solution, while this deviation is basically zero for the other reported wage rates.

the tax schedule's notch relevant. But even in this case the median solution deviates from the best solution by less than one tenth of one percent. Considering the top one percent of solutions reduces that deviation by one to two orders of magnitude, which justifies our adjusting GLO's algorithm to let it run to completion based on numerous initial guesses and then picking the best solution – a multi-start routine as is common in global optimization.

As an additional test for the GLO, we apply it to solving standard test functions for global optimization as Arnoud et al. (2019) use to solve up to 10-dimensional problems. We find that, as appendix C reports, the GLO can quickly, precisely, and robustly find the analytically known global optimum of these test functions in up to 100-dimensional problems.

4 Labor Supply Response to Kinks and Notches

In this section, we first analyse labor supply bunching as a response to convex kinks in the income-tax code leaving out FICA as well as the provision of basic income. Second, we add the FICA tax, which introduces a concave kink and thereby generates flipping behavior. Third, we examine the combination of flipping and bunching arising from including basic income and its associated notch. Finally, we discuss how kinks and notches can produce indeterminacy in the timing of labor supply, i.e., multiple choices of when to work, each of which delivers identical maximum lifetime utility. Most of this section assumes worker's wages are fixed in real terms throughout their working years.

4.1 Bunching Due to Convex Kinks From Income-Tax Brackets

We start by focusing on the two income-tax kinks that are most convex – the jump in the marginal tax rate from 12 to 22 percent occurring at $41,775 and the jump from 24 to 32 percent occurring at $170,050. Consider two workers earning wage rates of $35,000 and $40,000 for one unit of labor supply (i.e. one year of full-time work). As figure 6(a) shows, the lower-wage worker works substantially more than the higher-wage worker.[15] As detailed below, this reflects the high-wage worker's response to her much higher marginal tax rate. As a consequence, the higher-wage worker's consumption is only 7.5 percent higher while earning 14.3 percent more per unit of labor. Figure 6(b) provides a similar picture for the case of high-wage-rate households in the vicinity of the convex kink at $170,050.

Next, we consider the labor supply behavior for ranges of wage rates around the two highly convex kinks of the income tax code. Figure 7 shows how labor supply and labor income depend on the wage rate. For a substantial range of wage rates – from $37,500 to $39,500 – households chose to earn exactly $41,775, the amount corresponding to the major convex kink at the lower end of the tax schedule. Within this range of wage rates, labor supply is strongly decreasing in the pre-tax wage. Starting at a wage of about $40,000, households are willing to pay the higher marginal tax rate, and labor supply and income both increase. However, labor supply stays substantially below the labor supply of agents with lower wage rates. As shown in figure 8, the behavior is qualitatively similar at the other highly convex kink – $170,050. The range of bunching wage rates, however, is much larger, ranging between $154,000 and $163,000.

We now turn to bunching across time. Consider, in figure 9, bunching by a worker whose age-25 wage rate is $109,184 and rises by one percent annually.[16] This worker's labor supply remains low for several years. Then it begins to rise, but not by enough to push the worker above the 24 percent bracket. This continues until they reach the $170,050 bracket threshold. At this point, they reduce their labor supply each year for several years to avoid moving into the 32 percent bracket, which they eventually do.

[15] If we consider one unit of labor as the average number of hours worked in the U.S. in 2022, then a difference of five percentage points represents about 90 hours of work.

[16] Note that the present value of this wage path is equal to the present value of earning a constant annual wage of $130,000.

Figure 6: Life-Cycle Profiles of Households Facing Tax Code with Convex Kinks

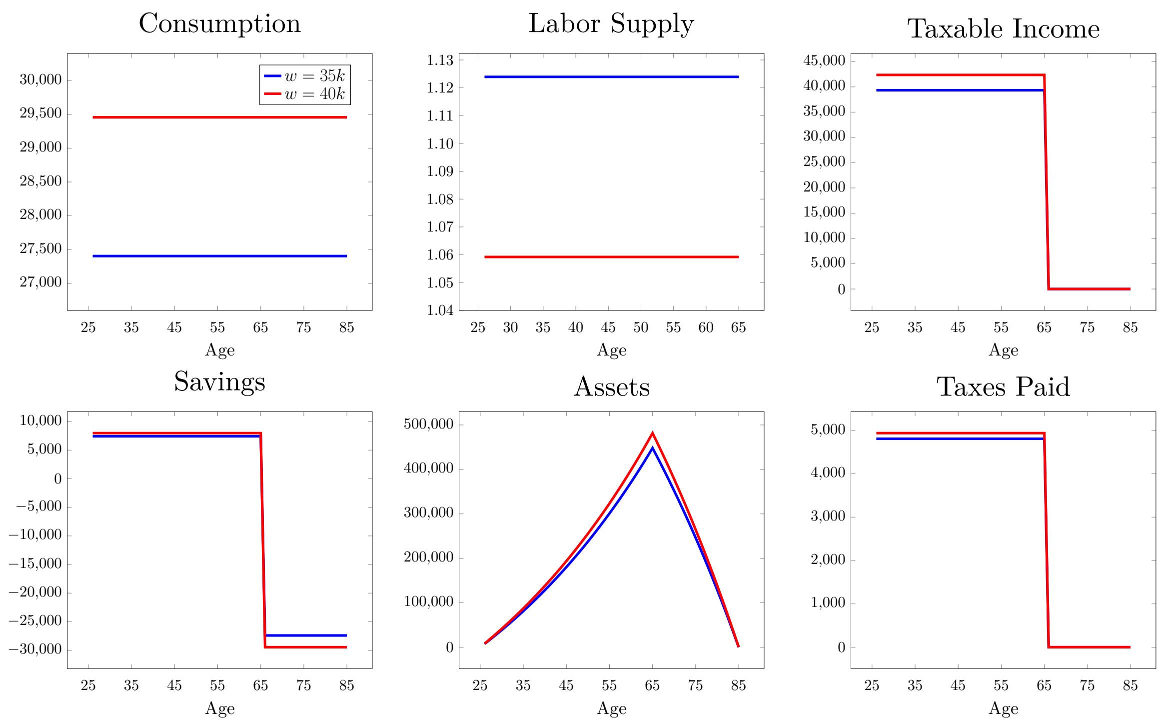

(a) Low-Wage-Rate Households

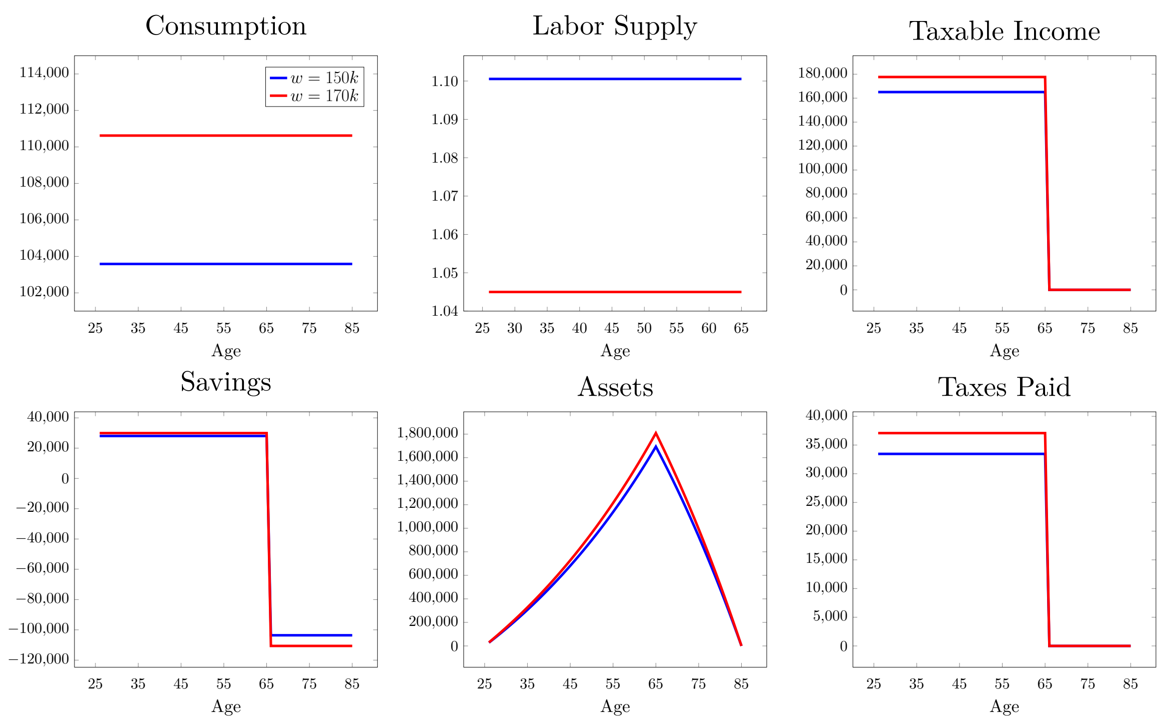

(b) High-Wage-Rate Households

Note: Life-cycle profiles for households facing labor taxation by the income tax code with convex kinks only (INC. tax). Consumption profiles are flat as interest rate and time preference rate are equal and asset income is not taxed. Both assumptions are for illustration and are relaxed below.

Figure 7: Bunching at the Convex Kink at $41,775

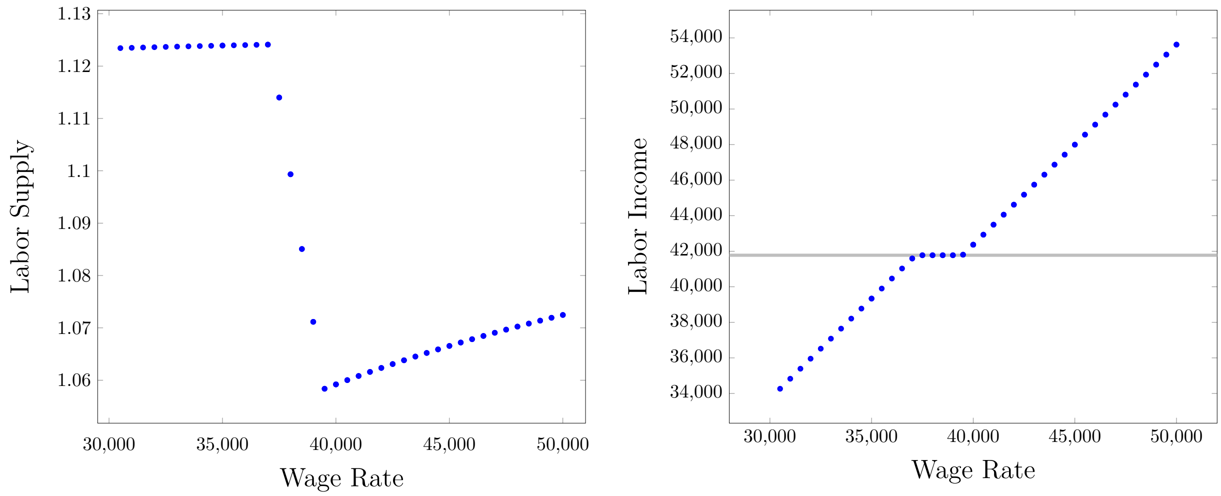

Note: Labor supply and labor income as a function of the wage rate. Bunching occurs on the grey line, which represents the boundary between two tax brackets.

Figure 8: Bunching at the Convex Kink at $170,050

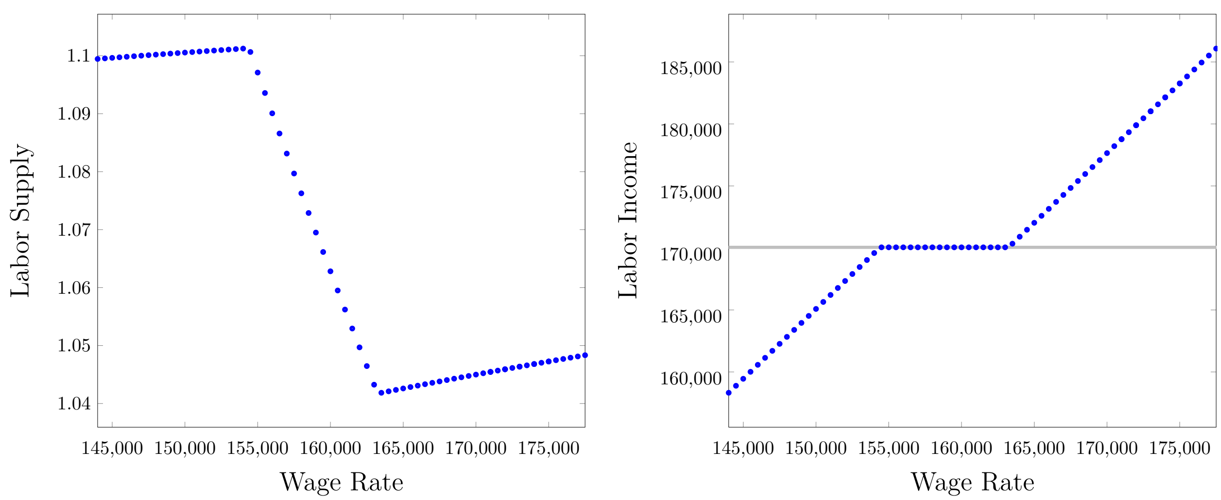

Note: Labor supply and labor income as a function of the wage rate. Bunching occurs on the grey line, which represents the boundary between two tax brackets.

Figure 9: Bunching over the Life-Cycle when the Wage Rate Grows

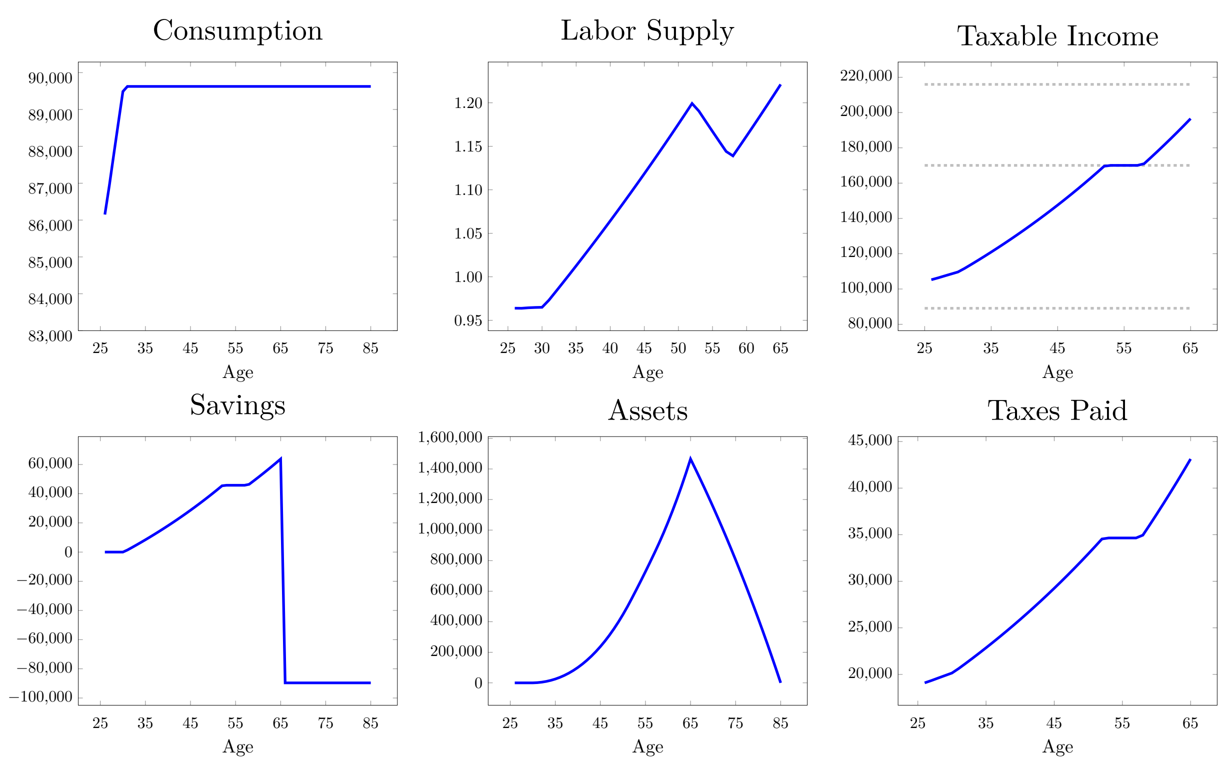

Note: Life-Cycle profiles of high-wage-rate household with one percent yearly wage growth.

Figure 10: Life-Cycle Profiles of Households Facing Income Plus FICA Tax

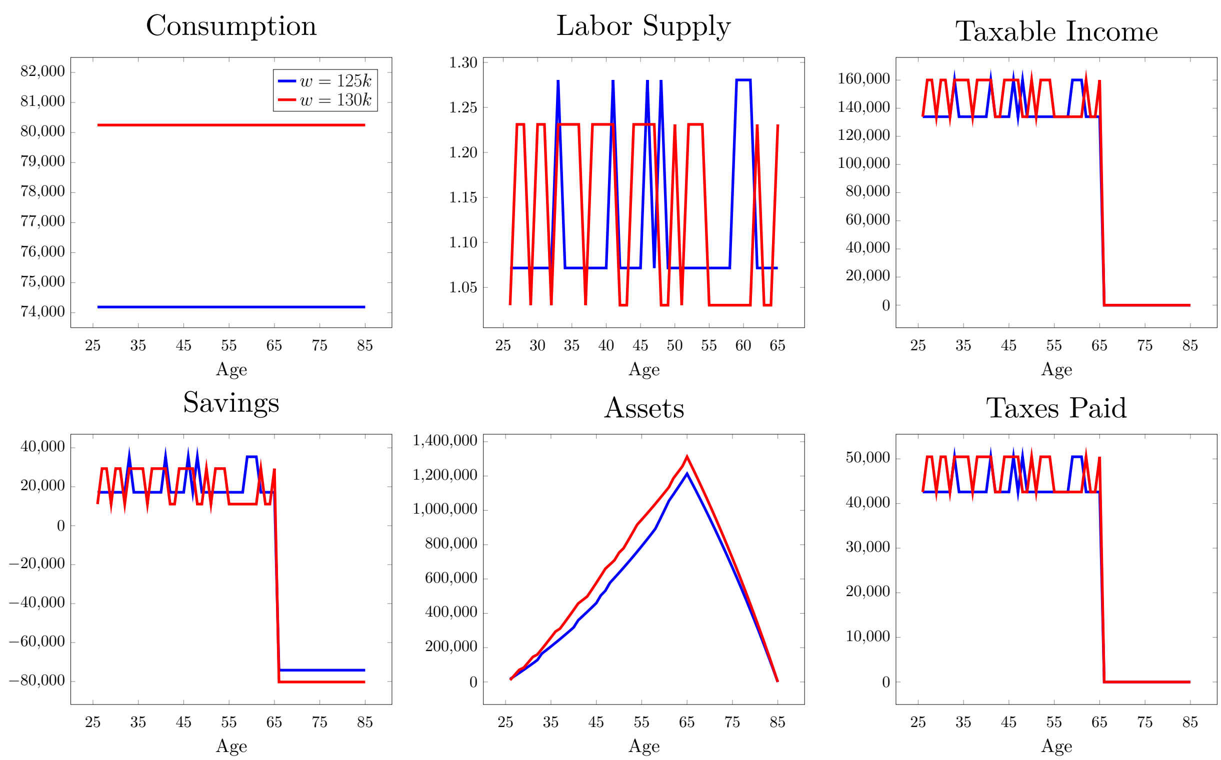

4.2 Flipping due to FICA's Concave Kink

We next include the 12.4 percent FICA payroll tax with, to repeat, its ceiling (concave kink) at \$147,000. Figure 10 displays life-cycle profiles for households with annual wage rates of \$125,000 and \$130,000. Despite having flat consumption profiles, their labor supply profiles exhibit frequent flipping. In particular, agents are jumping back and forth between two tax brackets – below and above the FICA threshold of \$147,000. While the lower bracket is characterized by a higher marginal tax rate, continually working sufficient hours of work to exceed the FICA threshold in all periods is sub-optimal. At the same time, it would also be sub-optimal to always stay below the threshold. Hence, the optimum entails working less in some years and more in others.

Figure 11 shows how labor supply and, thus, labor income are affected by the non-convex kink. Workers with wages between \$122,500 and \$135,500 flip. To be precise, they flip between low labor supply (the blue dots in the plot) that generates income substantially below the FICA threshold (grey line) and high labor supply (the red dots) that generates income much higher than the FICA threshold. Generating labor income close to the FICA threshold is never optimal – the exact opposite of bunching.

4.3 Bunching and Flipping Due to the Basic-Income Claw-Back Notch

We now turn to flipping and bunching due to the notch at \$15,000 in the tax and transfer scheme. Recall, the notch is generated by the assumption of a \$10,000 basic income that (single) households receive provided they earn less than \$15,000. Figure 12 displays the age-labor supply profile for

Figure 11: Flipping at the Concave Kink at $147,000

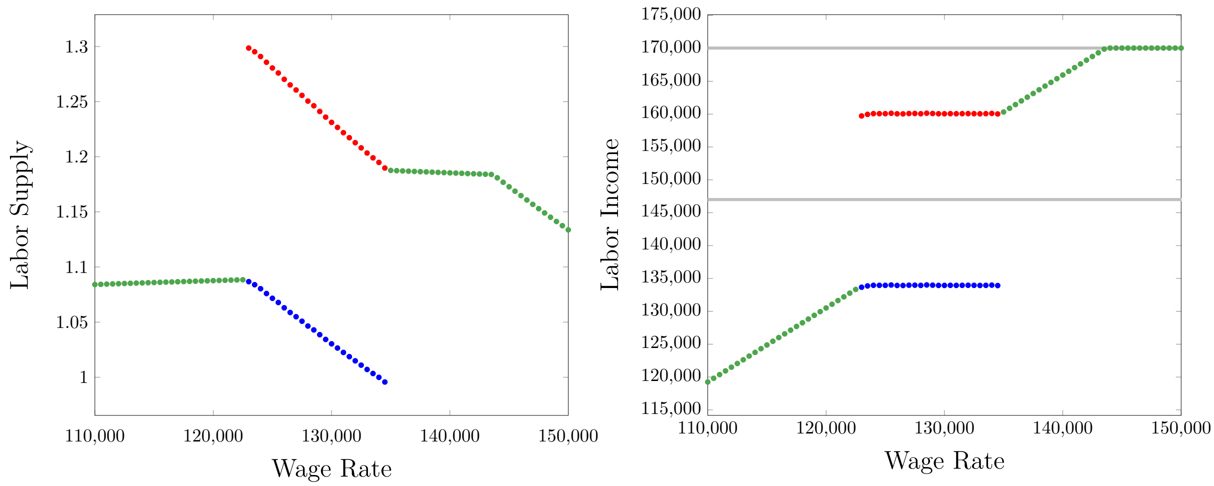

Note: Labor supply and labor income as a function of the wage rate. Flipping occurs in the neighborhood of the lower grey line that represents the FICA threshold. Households earning between $123,000 and $134,500 flip between the blue and the red points. Households with lower or higher wage rates do not flip, corresponding to the green points.

workers with three different wage rates – $30,000, $38,000, and $46,000. Figure 13 displays the labor supply and resulting labor income for a wide range of wage rates.

Workers with a wage rate below roughly $30,000 reduce their labor supply each year to stay below the $15,000 threshold. Starting from wage rates around $46,000 individuals choose to never collect the basic income. Between these values, however, it's optimal to collect basic income in some years and to earn substantially above the threshold in other years. In fact, it is never optimal to supply an amount of labor that generates anything strictly between the $15,000 threshold for the basic benefit and the upper bound for the 12 percent tax bracket at $41,775. Household who are flipping are saving substantial amounts in the periods they work. While one might argue that this is ruled out by assets tests in most real-world transfer schemes, we believe that such behavior is still relevant as there are often exceptions for assets like housing or cars and also because households can hide assets by 'gifting' to friends or relatives.

4.4 Annual Labor Supply – the Potential for Indeterminacy and Hypersensitivity

Depending on a worker's wage rate, the timing, if not the average level, of annual labor supply can be extremely sensitive to parameter values. Indeed, in the case we consider, where the interest rate and time preference rate are both precisely two percent, the timing of when flipping agents work part time and over time can be essentially indeterminate. This arises, in our model, around the FICA tax kink. Consider the $130,000 wage-rate worker facing the full tax system, where the concave kink induced by the FICA threshold induces flipping. Figure 14 displays consumption, labor supply, and assets over the life-cycle for three different solutions of the household's problem – the three best solutions from multiple different starts (different starting guesses for annual labor supply and consumption) of the GLO algorithm. The three solutions produce the same lifetime utility up to eight decimals. And

Figure 12: Life-Cycle Profiles of Low-Wage-Rate Households Facing Tax Code with Basic Income

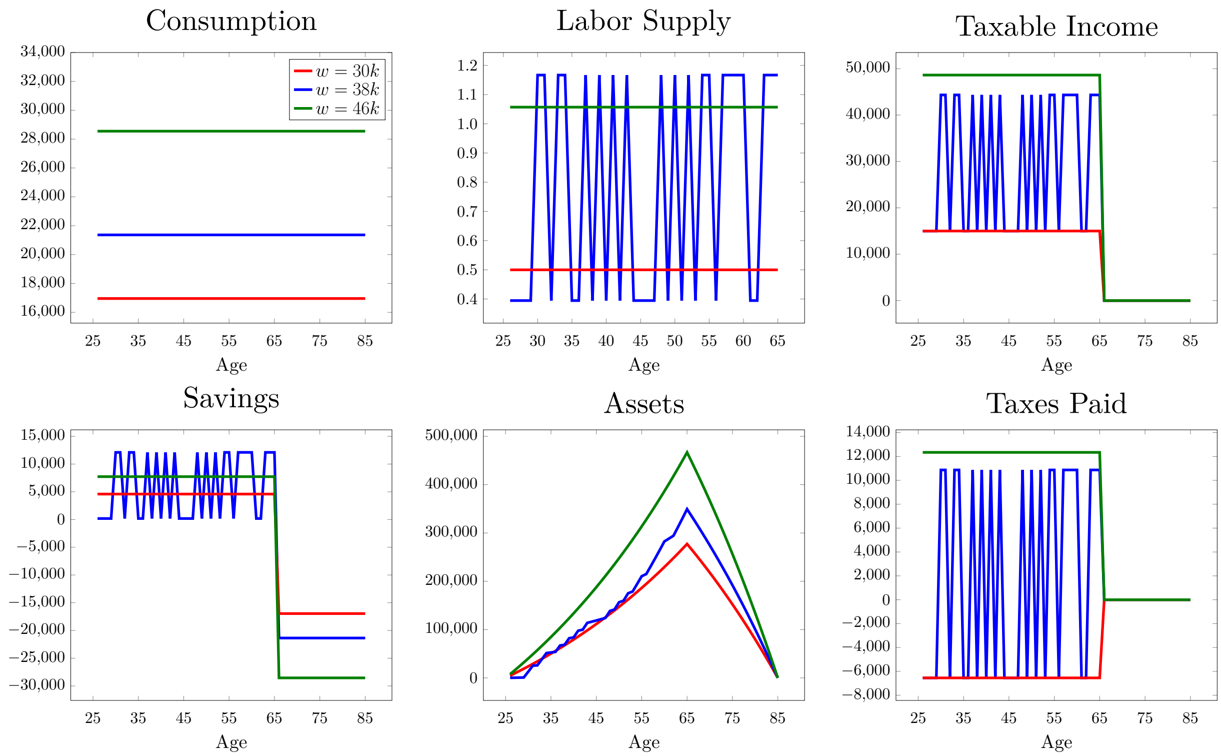

.

Note: Life-cycle profiles for three different wage rates, one resulting in below-threshold labor supply throughout working life, one in above-threshold labor supply, while the intermediate wage-rate implies flipping between staying below and above the threshold.

Figure 13: Bunching and Flipping at the Notch at $15,000

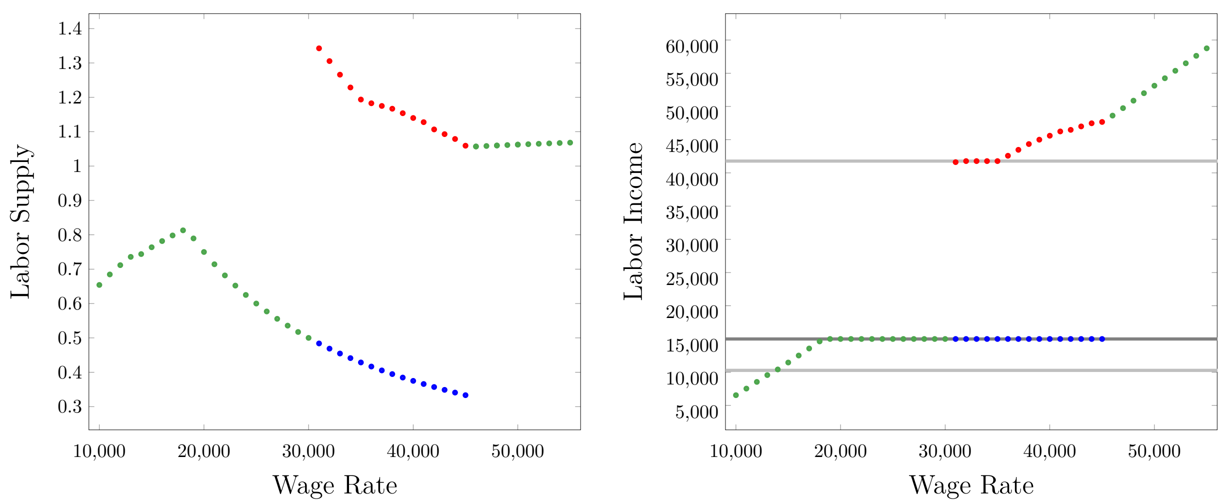

Note: Labor supply and labor income as a function of the wage rate. The notch at $15,000 corresponds to the lower grey line. For wage rates between $19,000 and $45,000 it induces bunching. In addition, for wage rates between $31,000 and $45,000, there is flipping between blue and red points.

Figure 14: Indeterminacy in Annual Labor Supply of Households Flipping Around FICA Threshold

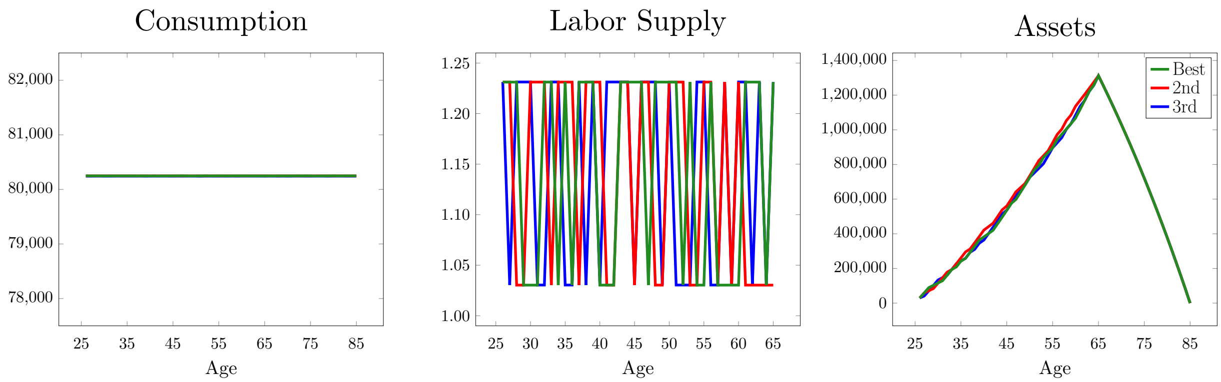

Note: Consumption, labor supply, and assets over the life-cycle for three different solutions of the problem of a household with wage rate $130,000. Lifetime utilities are identical up to the eighth decimal while the labor supply profiles differ strongly, in particular in terms of choice of periods in which the household works overtime (to exceed the FICA threshold) and in which it does not.

their age-consumption profiles (average consumption by age) differ by less than 0.01 percent. Yet, the timing of labor supply differs dramatically.[17]

Next, consider figure 15, which compares annual labor supply, saving, and consumption for two otherwise identical workers – one with a 2.0 percent and the other with a 1.9 percent time-preference rate. This minor difference in preferences produces a major difference in the age pattern of choice variables. When the worker's time-preference rate equals the interest rates, the age-consumption profile is, as expected, flat, but the annual earnings are indeterminate, flipping above and below the FICA threshold.[18] Yet, if the worker has a slightly lower time-preference rate, the age-consumption profile rises and the age-labor supply monotonically declines, dropping sharply when earnings fall below the FICA threshold.

[17] Interestingly, in the second solution, the household works overtime for 22 periods, but for 23 periods in the other two solutions. Hence, the household is basically indifferent not only about when to work, but also how much to work.

[18] Given separability in consumption and labor supply, the analytically optimal age-consumption profile is time-invariant absent binding cash-flow constraints.

Figure 15: Life-Cycle Profiles of $130,000 Worker with Different Time-Preference Rates

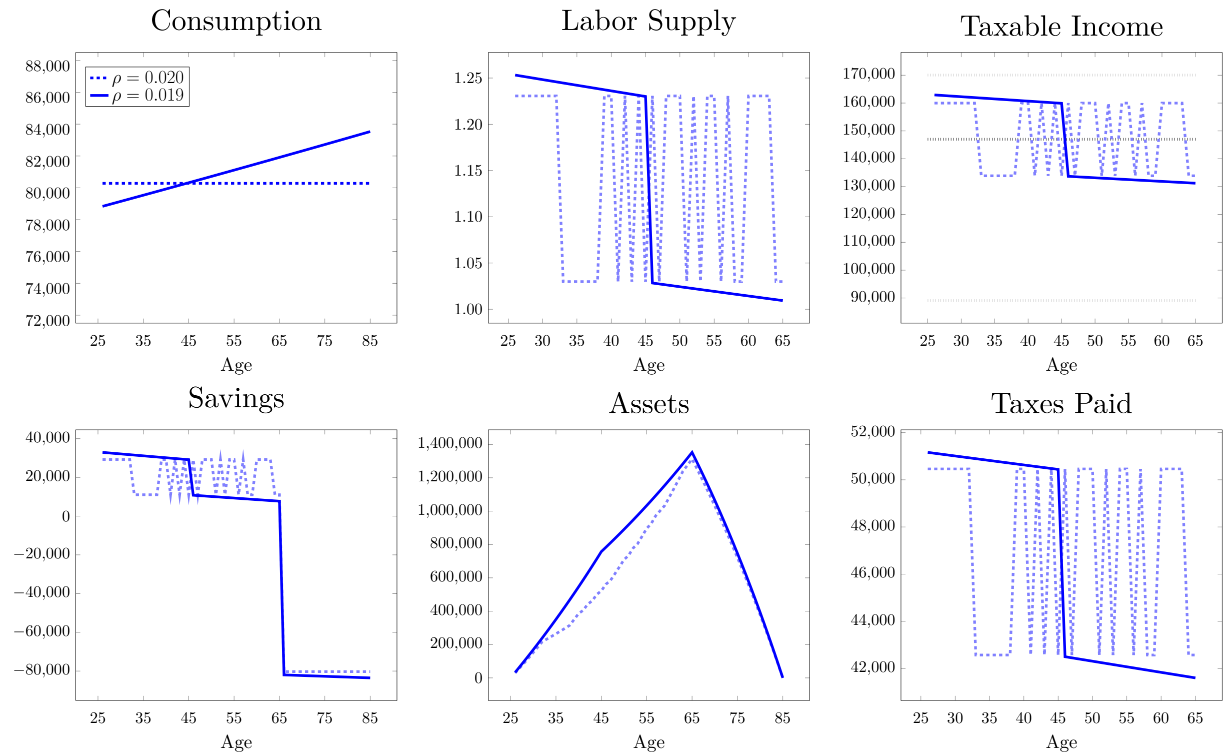

Note: Consumption, labor supply, and assets over the life-cycle for two different time-discount rates, ρ, while the interest rate in both cases is equal to 2 percent. The dotted gray line in the taxable-income panel indicates the FICA tax threshold.

5 Distortions of Labor Supply and Associated Excess Burdens

Table 4 reports our fiscal system's impact on labor supply and saving compared with that arising under lump-sum taxation. In considering the results, bear in mind that a simple linear tax would make no difference to labor supply. When households face just the income-tax schedule with its convex kinks, labor supply is reduced by five to sixteen percent and this distortion tends to increase in the wage rate.[19] When the FICA tax is added to the schedule, labor supply is distorted substantially more, especially for those with low wage rates. Interestingly, the labor distortion is now no longer monotonically increasing in the wage rate, as marginal tax rates are no longer monotone. Finally, when we add basic income and its clawback at $15,000, labor supply is massively distorted for low-wage households.

Of course, the labor supply distortions just described have a substantial impact on savings. Table 5 reports the percentage deviation of wealth at retirement relative to the corresponding lump-sum tax case. In our baseline tax scheme (last row) wealth accumulation is reduced substantially for all wage rates, with reductions ranging between 18% and 42%. Accumulated wealth inherits the non-monotonicity from the labor supply results. The largest reduction in savings is experienced by the low-wage rate workers, whereas the smallest reduction is experienced by households close to the FICA threshold.

[19] Intuitively, non-linearity accentuates substitution effects.

Table 4: Labor Supply Distortion

Ann. Wage	INC.	INC.+FICA	INC.+FICA+BASIC
25,000	−5.84	−11.51	−40.23
30,000	−5.84	−11.50	−51.14
35,000	−5.84	−11.50	−38.65
40,000	−11.04	−16.88	−28.95
50,000	−10.91	−16.65	−16.65
100,000	−11.69	−17.26	−17.26
130,000	−11.59	−13.79	−13.45
150,000	−11.55	−13.93	−13.93
200,000	−15.67	−16.41	−16.41

Notes: Percentage deviation of labor supply from the lump-sum taxation case for three different tax schemes (in rows) and many different wage rates (in columns).

Table 5: Impact on Wealth Accumulation

Ann. Wage	INC.	INC.+FICA	INC.+FICA+BASIC
25,000	−6.54	−14.56	−31.91
30,000	−6.54	−14.57	−42.16
35,000	−6.54	−14.58	−39.02
40,000	−12.32	−21.08	−30.78
50,000	−12.45	−21.29	−21.29
100,000	−13.94	−23.13	−23.13
130,000	−14.03	−18.31	−18.28
150,000	−14.07	−19.03	−21.93
200,000	−19.37	−22.24	−25.02

Notes: Percentage deviation of wealth at retirement from the lump-sum taxation case for three different tax schemes (in rows) and many different wage rates (in columns).

Analysis of the excess burden from taxation – the cost of distorting household consumption and leisure (labor supply and saving) decisions – dates to mid-19th century work by Jules Dupuit (see Ekelund Jr and Hébert, 2012). But Harberger (1964a) and Harberger (1964b) made excess burden analysis a mainstay of public finance. A multitude of studies applied Harberger's approach to all manner of tax-induced distortions. Notable examples include Browning (1975), who examined distortions in labor supply from payroll taxation, Feldstein (1978) who analyzed saving distortions from taxing capital income, and Rosen (1978) who considers the excess burden of wage taxation.[20][21] We build on these traditional studies by measuring excess burdens when fiscal systems are NND.

To measure the distortions produced by our tax system, we follow the standard procedure – compare, as a consumption equivalent variation (CEV), distorted lifetime utility with lifetime utility under lump-sum taxation. CEV tells us the percentage increase in annual consumption that will raise a distorted worker's lifetime utility to that enjoyed under lump-sum taxation. Appendix A provides the precise formula. To control for cash-flow constraints, we collect, from each worker at a given age,

[20] Auerbach (1985) and Auerbach and Hines (2002) provide extensive reviews of this literature.

[21] One shortcoming of Harberger's approach is the failure to include decision margins that are not directly affected by the specific tax under consideration. Lindsey (1987), Feldstein (1995) and Auten and Carroll (1999) addressed this issue by estimating the responsiveness of taxable income to changes in tax rates.

the same lump-sum net taxes as they pay when facing one of three distorted tax systems. The first includes just the federal income-tax kinks, the second includes those kinks, plus the concave kink from the FICA tax, and the third includes all kinks plus basic income and the notch from its clawback.

Table 6's third column shows that the full set of tax provisions produces huge distortions for low-wage workers. Take the $30,000 worker. Their labor supply and saving decisions are distorted by 26.3% of their post-tax living standard. Despite having a full-time wage that is twice as high as the basic income notch, they choose to earn only $15,000 per year for their entire working life. This leaves them paying negative net taxes of about $6,500 per year. Hence, the real tax burden this worker faces delivers no positive revenue. It simply reflects having their choice of labor supply and, thus saving, heavily distorted. Compared with lump-sum taxation, their annual labor supply is 51 percent lower and their wealth at retirement is 42 percent lower.

The massive fiscal burden arising solely from distorting behavior drops dramatically for workers earning $50,000 or more. But even an excess burden ranging from 3 to 5 percent is non-trivial. As the difference between the table's first two and third columns shows, federal income taxes produce, on their own, small excess burdens apart from those earning $200,000 pre-tax. It's moderate when we add in the concave FICA kink. The table's bottom line, then, is that providing basic income with a severe clawback may, inadvertently, lock the poor into poverty.

Table 6: Excess Burden of Different Tax Schemes

Ann. Wage	INC.	INC.+FICA	INC.+FICA+BASIC
25,000	0.42	2.04	13.82
30,000	0.42	2.04	26.32
35,000	0.42	2.04	26.08
40,000	1.52	4.47	16.60
50,000	1.54	4.52	4.52
100,000	1.91	5.31	5.31
130,000	1.92	3.56	3.56
150,000	1.93	3.46	3.46
200,000	3.74	4.85	4.85

Notes: Excess burden relative to lump-sum taxation for three different tax schemes (in rows) and many different wage rates (in columns). Excess burdens are measured by the consumption equivalent variation defined in Appendix A.

5.1 Sensitivity to Frisch Elasticity

Figure 7 examines labor supply and saving responses to our tax system, as well as their associated excess burdens, when the Frisch elasticity is 0.5 rather than 1. Labor supply is measured as the percentage change in its average value over the work span. The impact on saving is captured by the percentage change in wealth at retirement. Remarkably, the distortions in labor supply and wealth at retirement, measured relative to what would arise under lump-sum net taxation, remain remarkably large when the model's intrinsic force for substitution, the Frisch elasticity, is only half as big. Indeed, for workers with wage rates of $25,000 and $30,000, labor supply and saving responses are essentially unchanged. And excess burden is actually larger. The same is close to true for the worker in the table with a $130,000 wage rate. Although there is a somewhat smaller reduction in average labor supply, wealth at retirement actually falls by more with the lower elasticity. Even more surprising, the excess burden is somewhat higher. In the table's other six cases, the results are as expected – smaller behavioral distortions and excess burdens.

How can less willingness to change labor supply in response to a change in the net wage – a lower Frisch elasticity in the context of our preferences – lead to larger fiscal policy-induced behavioral changes and excess burdens? Our answer is that the Frisch elasticity impacts the choice of labor supply both with and without policy. A lower elasticity can position a worker near a notch or kink or far away depending on the nature of the policies adopted. Hence, a worker with a lower Frisch elasticity can face greater work disincentives than those with a higher Frisch elasticity because the labor and work distortions they face are greater. Stated differently, the Frisch elasticity tells us about responses to linear wage taxation, not responses to NND tax systems.

Table 7: Labor Supply, Saving, and Excess Burdens when Assuming a 50% Lower Frisch Elasticity

	Labor Supply		Wealth		Excess Burden	
Ann Wage	$\gamma = 1$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 0.5$
25,000	−40.23	−40.15	−31.91	−34.28	13.82	19.20
30,000	−51.14	−50.74	−42.16	−43.85	26.32	34.87
35,000	−38.65	−18.99	−39.02	−26.78	26.08	15.67
40,000	−28.95	−10.97	−30.78	−17.02	16.60	2.58
50,000	−16.65	−11.92	−21.29	−18.48	4.52	3.15
100,000	−17.26	−12.46	−23.13	−19.98	5.31	3.73
130,000	−13.45	−12.37	−18.28	−20.14	3.56	3.78
150,000	−13.93	−7.69	−21.93	−14.00	3.46	1.45
200,000	−16.41	−10.66	−25.02	−17.85	4.85	2.78

Notes: Percentage deviation of average labor supply and wealth at retirement in the full tax system from the respective lump-sum taxation case. Comparing our baseline case of $\gamma = 1$ with $\gamma = 0.5$. Excess burdens are measured by the consumption equivalent variation defined in Appendix A.

6 Extensions

Can the GLO handle discrete choice, joint taxation of spouses, and intertemporally intertwined tax systems, specifically, income, not wage, taxation? Yes, as this section shows. This is expected. Given GLO's solution method, there is nothing special about these cases or, indeed, incorporating additional elements, such as a distaste for participation in welfare programs or other fiscal provisions emphasized by Moffitt (1983).[22] Each of the extensions considered below use our original GLO parameters. Certainly, tweaking these parameters can improve performance. But the fact that tweaking isn't needed shows the robustness of the GLO algorithm.[23]

[22]Another example is incorporating adjustment costs from changing the amount of labor supplied from one year to the next. Such costs may be particularly important for employed workers.

[23]The GLO takes less than three minutes to process a household's globally optimal behavior. Hence, with parallel processing, one can, it seems, process an entire data set in several minutes. This remarkable speed can facilitate empirical analysis that incorporates benefit-program participation preferences, labor-supply adjustment costs, Akerlof (1978)'s efficient *tagging* of benefits to agent characteristics, and other factors.

6.1 Discrete Labor Supply

Figure 16 considers a $125,000-wage worker who is limited to working either full time or part time. Full time (half time) is defined as 1.2 (0.6) units of labor supply in a period. The left-hand panel is based on the same preferences used above. The right-hand panel multiplies the disutility of work in each period by 1.1. The solid blue curve shows the optimal age-labor supply profile when labor supply is discrete. The dotted curve shows the corresponding fully flexible labor-supply case. In the left panel, restricting labor supply leads to no flipping when flipping would otherwise arise. In the right panel, the opposite occurs: Optimal discrete choice features flipping when optimal flexible choice entails fixed annual labor supply. This has important implications for estimating structural labor-supply models. It suggests that assuming discrete choice to make one's model computationally tractable may be problematic.

Figure 16: Comparing Discrete and Continuous Labor Choice

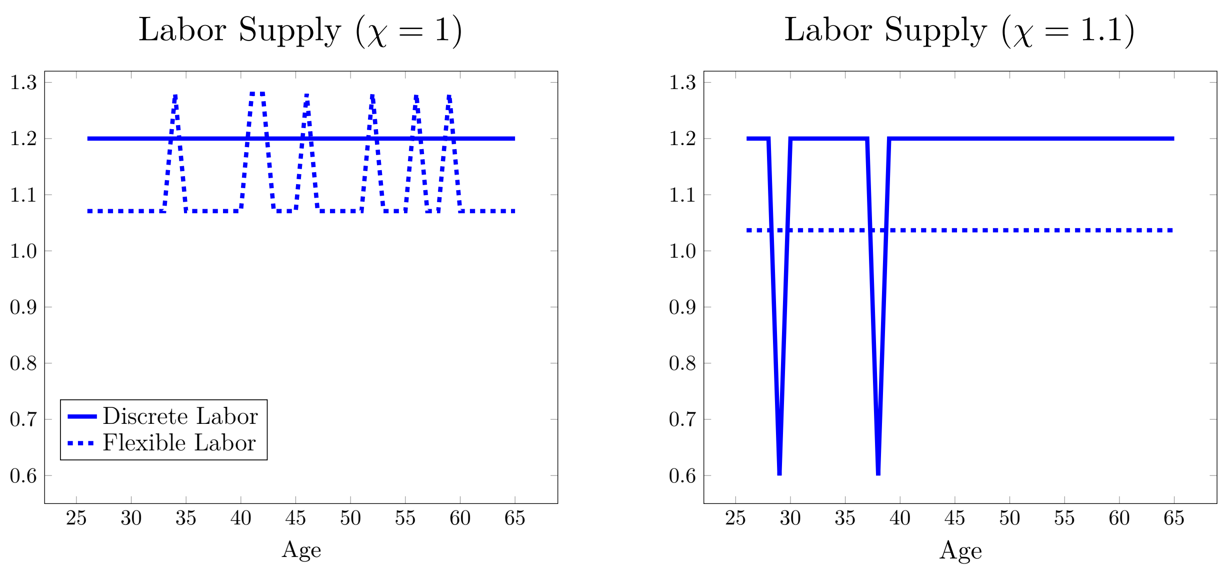

Note: Labor supply of worker with $125,000 wage rate who is either restricted to full-time or part-time work (discrete labor) or not restricted (flexible labor). The left-hand side assumes our standard disutility of labor parameter – $\chi = 1$. The right-hand assumes $\chi = 1.1$. In case of flexible labor, the household flips around the FICA threshold when $\chi = 1$, yet stays below that threshold when $\chi = 1.1$. In the discrete-labor case the household works full time throughout his working age when $\chi = 1$, yet occasionally works part time when $\chi = 1.1$.

6.2 Taxing Couples

Taxing couples on joint income renders the two spouses' labor-supply decisions interdependent. To examine this inter-dependency, we assume that both household members are the same age and live for T periods with joint lifetime utility given by

$$\sum_{t=1}^{T} \left(\frac{1}{1+\rho_t} \right)^{t-1} U(c_{1,t}, c_{2,t}, l_{1,t}, l_{2,t}), \tag{5}$$

Figure 17: Labor Supply when Couples are Taxed Jointly or Separately

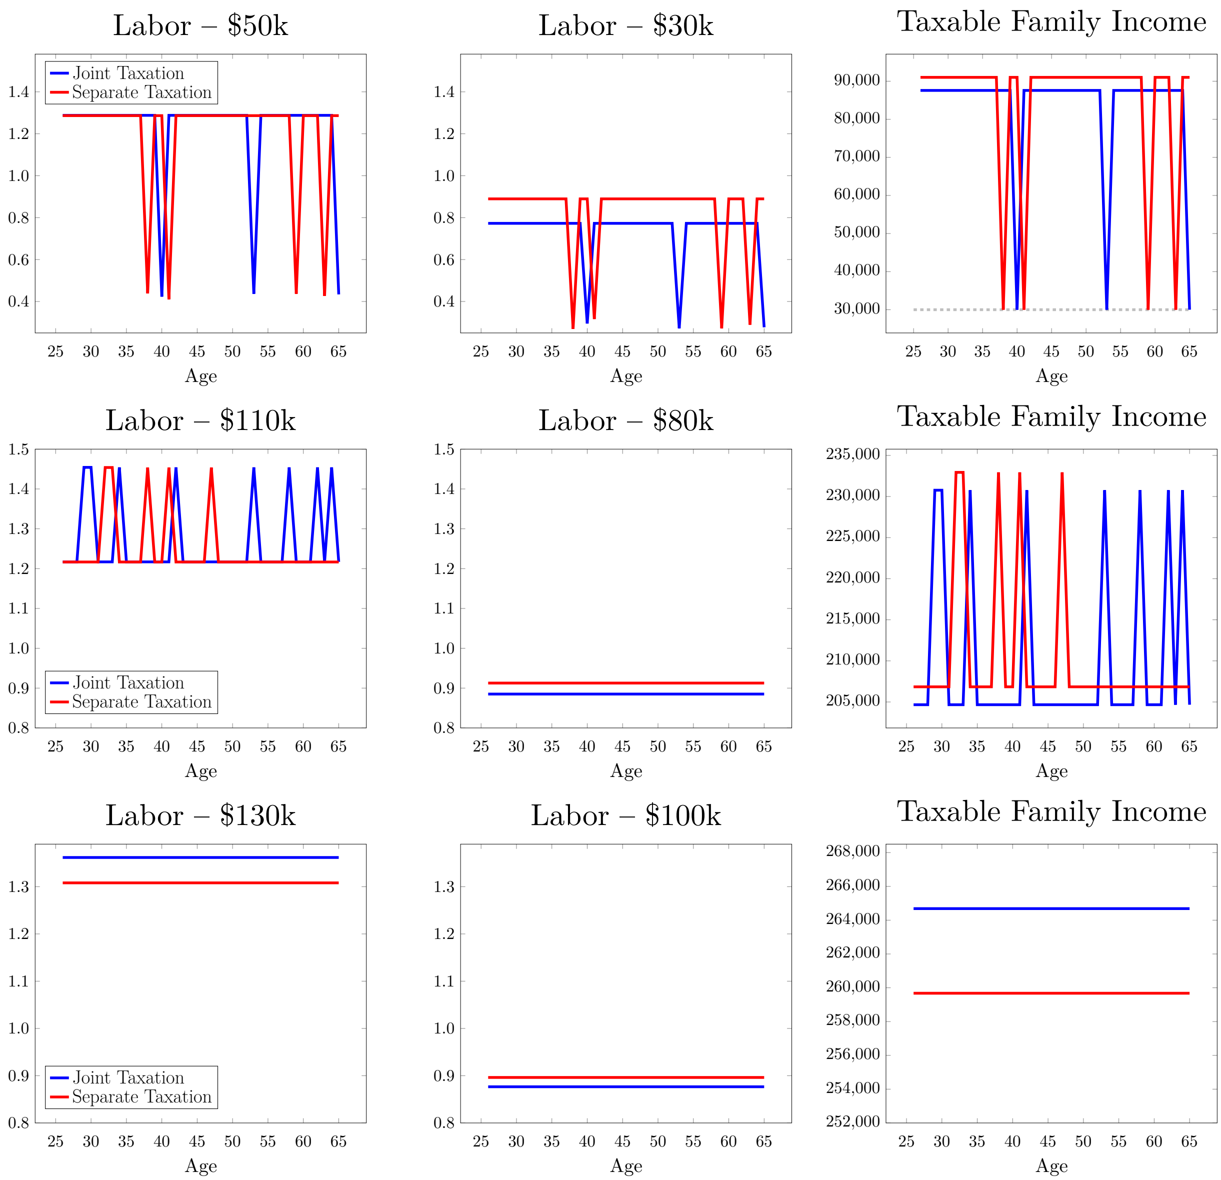

Note: Labor supply and taxable family income over the life cycle for three different types of couples. In the first case, both earners coordinate to collect basic income in some periods. In the second case, the primary earner flips around the FICA threshold. In the third case the primary earner increases labor supply substantially when taxed jointly to make use of lower marginal rates.

where c_i and l_i are consumption and labor supply of household member $i = 1, 2$. Per-period utility satisfies:

$$U(c_1, c_2, l_1, l_2) = 2 \cdot \log\left(\frac{c_1 + c_2}{2}\right) - \frac{l_1^{1+1/\gamma}}{1 + 1/\gamma} - \frac{l_2^{1+1/\gamma}}{1 + 1/\gamma}. \quad (6)$$

Thus, spouses value average consumption, whereas their two disutilities of labor supply are simply added together. Both spouses retire at 65, both have the same time-preference rate, and both have the same Frisch elasticity of labor supply. Although, positing couples facing joint taxation adds another labor-supply path over which the GLO must optimize, the program readily handles this. To be clear, we continue to randomly adjust two values of either consumption or labor supply in each iteration of our stochastic pattern-search routine. But now the two random adjustments are chosen from the set of annual labor supplies of each spouse and the annual levels of the two spouses' common consumption. Under *joint taxation*, household wage income is pooled and subject to the same marginal tax rates as in table 2, but with doubled levels of tax brackets. Under *separate taxation*, each spouse's wages are taxed separately according to table 2. The FICA tax treats spouses as single. As for the provision of basic income, couples whose joint labor earnings are less than $30,000 receive $20,000.

Figure 17 displays the labor supply of each spouse as well as taxable family labor income. Each row references a different couple with the higher earner's wage displayed in the first column and the lower earner's wage shown in the second column. The third column displays the couple's total taxable labor income. Blue lines reference the case of joint taxation and red lines denote separate taxation. The first row considers a relatively low-wage couple with wage rates of $50,000 and $30,000. There are two remarkable findings. First, under both tax schemes the two household members coordinate to work less in some periods in order to collect the basic income. Second, when comparing the two tax schedules, joint taxation discourages the secondary earner from working, leaving him working substantially fewer hours. This reduces overall taxable family income although joint taxation results in lower average taxes for any given pair of labor supply choices. This second observation also pertains to the case of our higher earning couple, with high and low spousal wage rates of $110,000 and $80,000, respectively. In this case, the primary earner flips her labor supply above and below the FICA threshold. The secondary earner, in contrast, works the same amount at all ages.

The third row considers a couple with even higher wage rates – $130,000 and $100,000. There is no longer any flipping: the primary earner stays above the FICA threshold while the secondary earner stays below. However, in this case, joint taxation not only lowers the secondary earner's labor supply. It also substantially increases the primary earner's labor supply. Consequently, taxable family income increases. In all three cases, the gap between the amount of labor supplied by the primary earner and the secondary earner increases in switching from separate to joint taxation. Joint taxation is always preferred by our married households. The welfare gain is substantial for the low-income case – roughly a 5 percent gain. It's marginal for the higher-earning couple – less than one percent.

6.3 Taxing Total Income

So far, we've focused on labor-income taxation to isolate its impact on labor supply. We now apply the federal tax schedule and the basic income clawback thereshold to total income, while the FICA taxation still applies just to labor income. Figure 18 displays age-labor supply profiles for workers with wage rates of $30,000, $38,000, and $46,000 as in figure 12. The consumption profiles are now generally downward sloping as capital-income taxation discourages saving. Stated differently, it encourages current over future consumption. Similarly labor supply tends to be upward sloping since capital-income taxation favors working more when old (taking leisure earlier). However, labor supply exhibits a much richer behavior. The red line shows slightly downward sloping labor supply as long as the household collects basic income. The blue line exhibits flipping for the first half of the working life and then increasing work. The green line flips down once and exhibits upward sloping behavior for the rest of the agent's pre-retirement years.

Figure 18: Life-Cycle Profiles when Taxing Total Income

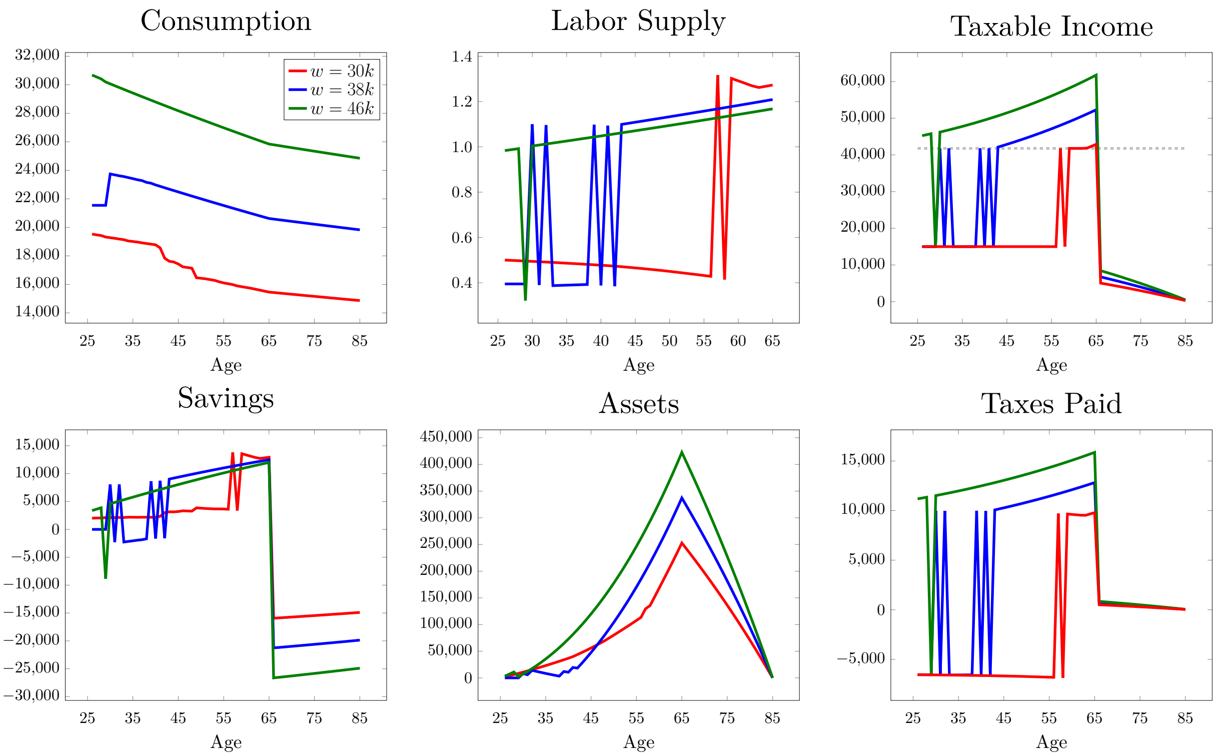

7 Conclusion

This paper develops the Global Life-Cycle Optimizer to study the impacts typical elements of fiscal policy can have on work and saving decisions. The GLO is a stochastic pattern-search algorithm specifically designed to determine optimal economic behavior in the context of non-differentiable, non-convex, and discontinuous (NND) choice sets. Economists have long recognized the NND nature of household budgets and the need for global optimization. But computational and algorithmic limitations generally restricted their analyses to static models with a limited set of policies, labor-supply options, or both. The GLO changes this playing field.

To demonstrate the GLO's potential, we consider labor income as well as total income taxation within a simplified U.S. fiscal system comprising three elements – the federal personal income-tax brackets, the Social Security payroll tax with its ceiling, and the provision of basic income to workers earning below a threshold. This tax system comprises eight differently-sized kinks (seven convex and one concave) and one notch.

The GLO readily reproduces the anomalous behavior that theory predicts and data record – the bunching of earnings just below benefit-eligibility or higher marginal-tax thresholds. But it also produces a) flipping – switching in different years between over-time work (high labor supply) and part-time work (low labor supply), b) rising, then falling, then rising labor supply over the work span of workers with increasing real wages seeking to avoid bracket creep, and c) potential extreme sensitivity of the level and pattern of annual labor supply to parameter values. Of most surprise is the wide range of wages over which labor supply is dramatically altered. This is particularly true for low-wage workers, some of who reduce their labor supply by more than 50 percent in the course of bunching their wages. Compared with lump-sum taxation, low-wage workers with a unitary Frisch elasticity of labor supply, experience excess burdens exceeding 26 percent of lifetime spending. For those unaffected by the notch, excess burdens are moderate.

As we show, the GLO can handle discrete choice, joint taxation of married couple's labor supplies, and taxation based on alternative tax bases. Furthermore, the GLO seems fully capable of simultaneously handling cash-flow constraints, fixed costs of working, minimum and maximum hours restrictions, labor supply adjustment costs, and benefit-program participation costs.[24] Inclusion of adjustment costs would surely limit the amount of flipping.[25] Agents who would otherwise flip absent adjustment costs would likely choose to work full time for part of their work span and part time for the remainder. This would manifest as early retirement. As for empirical analysis, preferences can, as in Burtless and Hausman (1978), be unobserved by the econometrician and, as in MaCurdy et al. (1990), variables can be measured with error. Simulated maximum likelihood (Lerman and Manski, 1981) as well as simulated method of moments (McFadden, 1989; Pakes and Pollard, 1989) can be used in conjunction with the GLO to measure preference and other parameters, such as adjustment costs governing switching from part time to over time work. The GLO also seems ideal for evaluating tax reforms. With a large and extensive data set and careful delineation of prevailing fiscal policies, one can compare GLO's results under current policy with that under proposed alternative policies – and do so on a state-by-state basis. This can include radical tax simplifications that can be predicated, as in Akerlof (1978), on observable characteristics. Comparisons can also cover macro aggregates as well as household-by-household labor supply, saving, and welfare changes. Hence, the tool can be used to search for Pareto-improving reforms. Although the GLO's capacities to handle the full U.S. fiscal system remain to be seen, preliminary work with the Fiscal Analyzer, which encompasses the entire potpourri of U.S. federal and state policy (see Auerbach et al., 2023), suggests no difficulties. Intuitively, altering GLO's annual net tax function changes nothing fundamental.

A tougher challenge is incorporating uncertainty beyond that facing econometricians, namely macro shocks and idiosyncratic uncertainties impacting households. One approach, which could readily be implemented with our current version of the GLO, is to follow Cai and Judd (2023) and assume certainty-equivalent behavior. Agents, in this case, make decisions in each period as if all future shocks equal their expected values. This transforms lifetime uncertainty problems into sequences of deterministic problems, which, as we've seen, the GLO can readily handle. The concern here is whether this approach properly captures risk averse behavior, specifically precautionary consumption and labor supply. The standard, but far more challenging path toward handling uncertainty requires the use of dynamic programming. A variant of the GLO might be used to find optimal policies in each step of value function iteration. While the GLO will likely be up to this task, a major additional challenge arises: interpolating the resulting non-differentiable value function accurately.[26] A more straightforward extension of our approach, and one on which we are embarked, is to simply specify the GLO's objective function as remaining lifetime expected utility and determine optimal decisions in NND settings along any sequence of future random outcome paths in the context of making globally optimal decisions along all other such paths. Conceptually speaking, this expanded program is fundamentally identical to the one we've solved. However, it is a far higher-dimensional problem than the one studied here, thus requiring practical limits on the degrees and types of uncertainty that can be incorporated – in particular limiting the number of periods in which shocks can occur.

A final point. One might reasonably ask whether individual households can make the calculations being modeled and processed by GLO. Complex fiscal provisions are, after all, a postwar phenomenon,

[24] Moffitt (1983), Fraker and Moffitt (1988), and Moffitt (1992) incorporate decisions over plan participation.

[25] Adjustment costs may reflect search time needed to find jobs that permit desired hours of work. They can also reflect loss of firm-specific human capital, reduced accumulation of human capital, and depreciation of human capital.

[26] A potential way to interpolate the value function in several dimensions while preserving monotonicity and concavity is Delaunay interpolation as employed in Brumm and Grill (2014).

not problems human brains have evolved to solve. Our response is twofold. First, observed earnings bunching shows that households can comprehend and appropriately respond to at least some forms of fiscal non-linearities. Second, over time, computation algorithms, like the GLO, will surely assist households in making optimal life-cycle labor supply and saving decisions. Future economists will likely prescribe optimal life-cycle choices not study computational mistakes that are easy to mischaracterize as behavioral shortcomings.

APPENDIX

A Consumption Equivalent Variation

Denote lifetime utility by

$$W(c,l) = \sum_{t=1}^{T} \beta^{t-1} U(c_t, l_t), \text{ with } \beta = \left(\frac{1}{1+\rho}\right). \tag{7}$$

Let (c^0, l^0) and (c^1, l^1) be consumption and labor supply under the considered tax system and under lump-sum taxation, respectively. Then the consumption equivalent variation (CEV) is defined by the following equality:

$$W(c^0(1+CEV), l^0) = W(c^1, l^1).$$

For the LHS we get:

$$\begin{aligned}
W(c^0(1+CEV), l^0) &= \sum_{t=1}^{T} \beta^{t-1} U(c_t^0(1+CEV), l_t^0) \\
&= \sum_{t=1}^{T} \beta^{t-1} \left(\log[c_t^0(1+CEV)] - \chi \frac{(l_t^0)^{1+1/\gamma}}{1+1/\gamma} \right) \\
&= \sum_{t=1}^{T} \beta^{t-1} \left(\log(c_t^0) + \log(1+CEV) - \chi \frac{(l_t^0)^{1+1/\gamma}}{1+1/\gamma} \right) \\
&= \sum_{t=1}^{T} \beta^{t-1} \left(\log(c_t^0) - \chi \frac{(l_t^0)^{1+1/\gamma}}{1+1/\gamma} \right) + \sum_{t=1}^{T} \beta^{t-1} \log(1+CEV) \\
&= W(c^0, l^0) + \sum_{t=1}^{T} \beta^{t-1} \log(1+CEV).
\end{aligned}$$

Combining the above equations results in:

$$\begin{aligned}
& W(c^0(1+CEV), l^0) = W(c^1, l^1) \\
\Leftrightarrow\ & W(c^0, l^0) + \sum_{t=1}^{T} \beta^{t-1} \log(1+CEV) = W(c^1, l^1) \\
\Leftrightarrow\ & \sum_{t=1}^{T} \beta^{t-1} \log(1+CEV) = W(c^1, l^1) - W(c^0, l^0) \\
\Leftrightarrow\ & CEV = \exp\left(\frac{W(c^1, l^1) - W(c^0, l^0)}{\sum_{t=1}^{T} \beta^{t-1}} \right) - 1.
\end{aligned}$$

B GLO Algorithm

This appendix provides a formal description of the working of GLO for a single starting guess. We avoid the outer loop for multiple starting guesses to simplify notation and as its implementation is trivial: a final step simply chooses the solution with the highest objective value among all solution resulting from different starting guesses. For the above results, we ran the GLO with the following parameter values: $m^0 = 1$, $J = 500$, $\lambda = 0.0001$, $b_l = 0$, $b_u = 1.25$, $\varepsilon = 10^{-8}$.

1. **Initialization:** Set initial path for consumption and gross labor income

$$x^0 = \{c_1^0, \ldots, c_T^0, w_1 l_1^0, \ldots, w_R l_R^0\}$$

of size $N = T + R$, stopping criterion $\varepsilon > 0$, initial mesh size $m^0 > \varepsilon$, adjustment interval $[b_l, b_u]$, and penalty parameter $\lambda > 0$ for the objective function

$$W_\lambda(x) = \sum_{t=1}^{T} \left(\frac{1}{1+\rho_t}\right)^{t-1} U(c_t, l_t) - \sum_{t=1}^{T-1} \lambda \min\{0, a_{t+1}\}^2,$$

where $l_t = x_{T+t}/w_t$ and $a_{t+1} = w_t l_t + (1 + r_t)a_t - \mathcal{T}(w_t l_t) - c_t$ for $t = 1, \ldots, T$.

2. **Construction of poll set:** Given the current point x^k and mesh size m^k the poll set contains J points of dimension N,

$$\mathcal{P} = \{p_1, \ldots, p_J\}.$$

To generate p_j, $j = 1, \ldots, J$, execute the following steps:

(a) Draw two distinct random integers $j_1, j_2 \in \{1, \ldots, N\}$ and one uniformly distributed random real number $z \in [b_l, b_u]$.

(b) Compute four different candidates $\tilde{p}_{i,1}, \ldots, \tilde{p}_{i,4}$, which are identical to x^k except for entries j_1 and j_2. In particular, set $\tilde{p}_{i,1} = \cdots = \tilde{p}_{i,4} = x^k$ and add and subtract m^k to j_1, and $z \cdot m^k$ to j_2:

$$\begin{aligned}
\tilde{p}_{i,1}(j_1) &= x^k(j_1) + m, \quad & \tilde{p}_{i,1}(j_2) &= x^k(j_2) + z \cdot m \\
\tilde{p}_{i,2}(j_1) &= x^k(j_1) + m, \quad & \tilde{p}_{i,2}(j_2) &= x^k(j_2) - z \cdot m \\
\tilde{p}_{i,3}(j_1) &= x^k(j_1) - m, \quad & \tilde{p}_{i,3}(j_2) &= x^k(j_2) + z \cdot m \\
\tilde{p}_{i,4}(j_1) &= x^k(j_1) - m, \quad & \tilde{p}_{i,4}(j_2) &= x^k(j_2) - z \cdot m
\end{aligned}$$

(c) Adjust terminal consumption levels $\tilde{p}_{i,T}$, $i = 1, \ldots, 4$, such that the lifetime budget constraint is satisfied: $c_T = (1 + r_T)a_T$, where a_T is calculated as in step 1.

(d) Choose p_i to be the candidate with the highest objective value:

$$p_i = \underset{\{\tilde{p}_{i,1}, \ldots, \tilde{p}_{i,4}\}}{\arg\max} \; W_\lambda(p).$$

3. **Evaluation of poll set**: Pick the p_i with the highest objective value:

$$p^* = \arg\max_{\mathcal{P}} W_\lambda(p).$$

If $W_\lambda(p^*) > W_\lambda(x^k)$, then the poll is successful: set $x^{k+1} = p^*$ and $m^{k+1} = m^k \cdot 2$.
If $W_\lambda(p^*) \leq W_\lambda(x^k)$, the poll was unsuccessful: set $x^{k+1} = x^k$ and $m^{k+1} = m^k/2$.

4. **Convergence**: If $m^k < \varepsilon$, stop and return p^* as solution; otherwise, go to step 2.

C Global Optimization of Standard Test Functions

We follow Arnoud et al. (2019) and use standard test functions for global optimizers to benchmark the performance of GLO. In particular, we use the following three test functions:

- Levi function No. 13:

$$f(x) = \sin^2(3\pi x_1) + (x_n - 1)^2[1 + \sin^2(2\pi x_n)] + \sum_{i=1}^{n-1}(x_i - 1)^2[1 + \sin^2(3\pi x_{i+1})] + 1, \quad (8)$$

with $x \in [-10, 10]^n$, which has its global minimum at $x = (1, \ldots, 1)$ with function value $f(1, \ldots, 1) = 1$.

- Rastrigin function:

$$f(x) = An + \sum_{i=1}^{n}[x_i^2 - A\cos(2\pi x_i)] + 1, \quad (9)$$

with $A = 10$ and $x \in [-5.12, 5.12]^n$, which has its global minimum at $x = (0, \ldots, 0)$ with function value $f(0, \ldots, 0) = 1$.

- Griewank function:

$$f(x) = \sum_{i=1}^{n} \frac{x_i^2}{a} - \prod_{i=1}^{n} \cos\left(\frac{x_i}{\sqrt{i}}\right) + 2, \quad (10)$$

with $a = 200$ and $x \in [-100, 100]^n$, which has its global minimum at $x = (0, \ldots, 0)$ with function value $f(0, \ldots, 0) = 1$.

Besides the change of objective function, we only make two obvious and two additional minor adjustments to the GLO. Clearly, we are now minimizing instead of maximizing and step 2 (c) is obsolete as we are now solving unconstrained optimization problems. Moreover, we choose $J = 2000$, i.e., we increase the size of the poll set in each iteration. This choice guarantees a success rate of at least 50% in all cases of the test functions as reported below. Finally, we set $\varepsilon = 10^{-6}$. This choice does not affect the success rate, as GLO either finds the global minimum or gets stuck in a local one, but it lowers the compute time to some degree.

Figure 19 displays the performance of the GLO on the three considered test functions of varying dimensionality. The left panel displays the compute time for 100 GLO runs, which turns out to grow much less than exponentially (concave shape with logarithmic time scale). The GLO can swiftly solve all three test functions in up to 100 dimensions.[27] It does so at a remarkable success rate of more than 50% across all test functions and dimensions considered, as the right panel of figure 19 shows the success rates of GLO to find the respective global minimum. We define a trial as successful if

[27] Note that Arnoud et al. (2019) reports results for up to ten dimensional problems only. Note also that that paper uses a fourth test function, the Rosenbrock function. As the GLO can only solve that function for up to about 12 dimensions, we do not include the results here.

the absolute distance between the objective function value found by GLO and the global optimum is less than 10^{-6}. While GLO finds the global optimum for the Levi No. 13 function in all trials, the success rate is increasing in the number of dimensions for the Griewank function and decreasing for the Rastrigin function. Note that that these success rates of the GLO for single starting values translate to very high success rates for the GLO with multiple starting guesses. Even in the worst case, the 100-dimensional Rastrigin function, an implementation with 10 starting guesses has a success rate above 99.9%.

Figure 19: GLO's Performance on Standard Test Functions

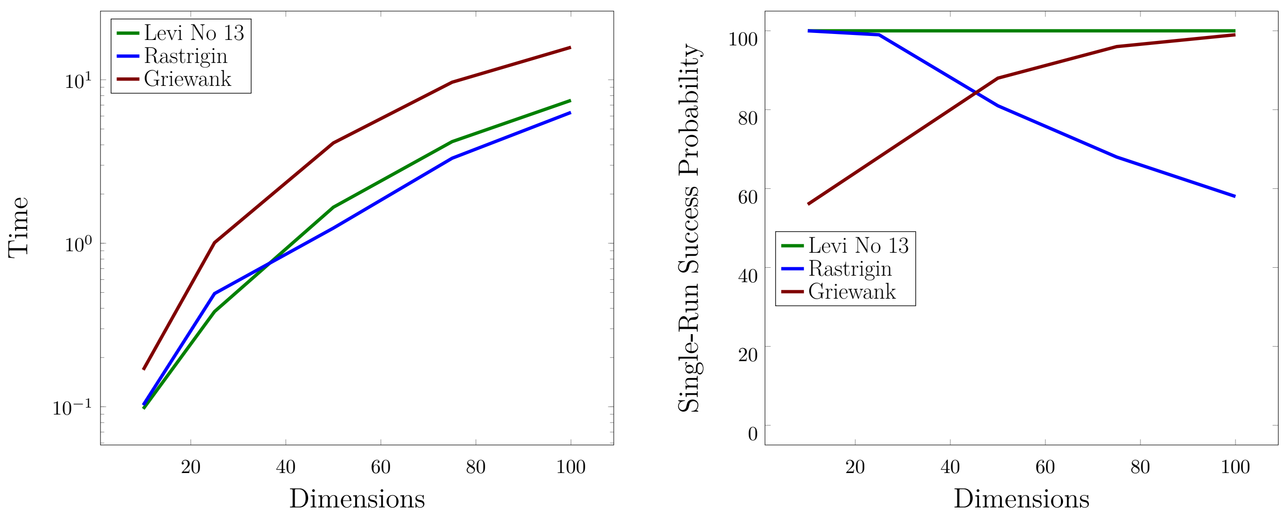

Note: The left panel displays the average compute time (in seconds) for a single run of the GLO when solving the respective test functions. The right panel shows the success probability of the GLO for a single starting guess in finding the analytically known global minima of these test functions. With 10 starting guesses, success rates are above 99.9% for all considered specifications.

References

Akerlof, G. A. (1978). The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *American Economic Review*, 68(1):8–19.

Arnoud, A., Guvenen, F., and Kleineberg, T. (2019). Benchmarking global optimizers. NBER Working Paper 26340.

Audet, C. and Dennis, J. (2003). Analysis of generalized pattern searches. *SIAM Journal on Optimization*, 13(3):889–903.

Auerbach, A. J. (1985). The theory of excess burden and optimal taxation. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 1, pages 61–127. Elsevier.

Auerbach, A. J. and Hines, J. R. (2002). Taxation and economic efficiency. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 3, pages 1347–1421. Elsevier.

Auerbach, A. J., Kotlikoff, L. J., and Koehler, D. (2023). U.S. inequality and fiscal progressivity: An intragenerational accounting. *Journal of Political Economy*, 131(5):1249–1293.

Auten, G. and Carroll, R. (1999). The effect of income taxes on household income. *Review of Economics and Statistics*, 81(4):681–693.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

Bastani, S. and Selin, H. (2014). Bunching and non-bunching at kink points of the swedish tax schedule. *Journal of Public Economics*, 109:36–49.

Bertanha, M., McCallum, A. H., and Seegert, N. (2023). Better bunching, nicer notching. *Journal of Econometrics*, 237(2):105512.

Bick, A. and Fuchs-Schündeln, N. (2018). Taxation and labour supply of married couples across countries: A macroeconomic analysis. *The Review of Economic Studies*, 85(3):1543–1576.

Blomquist, S., Newey, W. K., Kumar, A., and Liang, C.-Y. (2021). On bunching and identification of the taxable income elasticity. *Journal of Political Economy*, 129(8):2320–2343.

Blundell, R. and Hoynes, H. W. (2004). Has' in-work'benefit reform helped the labor market? In *Seeking a Premier Economy: The Economic Effects of British Economic Reforms, 1980-2000*, pages 411–460. University of Chicago Press.

Blundell, R. and MaCurdy, T. (1999). Labor supply: A review of alternative approaches. In Ashenfelter, O. C. and Card, D., editors, *Handbook of Labor Economics*, volume 3, pages 1559–1695. Elsevier.

Blundell, R. and Shephard, A. (2012). Employment, hours of work and the optimal taxation of low-income families. *The Review of Economic Studies*, 79(2):481–510.

Brown, K. M. (2013). The link between pensions and retirement timing: Lessons from California teachers. *Journal of Public Economics*, 98:1–14.

Browning, E. K. (1975). Labor supply distortions of social security. *Southern Economic Journal*, 42(2):243–252.

Brumm, J. and Grill, M. (2014). Computing equilibria in dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control*, 38:142–160.

Brumm, J. and Scheidegger, S. (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612.

Burtless, G. (1976). Taxes, transfers and preferences for work among black married men. Mimeographed. Cambridge, Mass.: Massachusetts Inst. Tech.

Burtless, G. and Hausman, J. A. (1978). The effect of taxation on labor supply: Evaluating the gary negative income tax experiment. *Journal of Political Economy*, 86(6):1103–1130.

Cai, Y. and Judd, K. L. (2023). A simple but powerful simulated certainty equivalent approximation method for dynamic stochastic problems. *Quantitative Economics*, 14(2):651–687.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2011a). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. *The Quarterly Journal of Economics*, 126(2):749–804.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011b). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–475.

Danzinger, S. and Plotnick, R. (1981). Chapter 6: Income maintenance programs and the pursuit of income security. *The ANNALS of the American Academy of Political and Social Science*, 453(1):130–152.

Duarte, V., Fonseca, J., Goodman, A. S., and Parker, J. A. (2021). Simple allocation rules and optimal portfolio choice over the lifecycle. NBER Working Paper 29559.

Ekelund Jr, R. B. and Hébert, R. F. (2012). The intellectual legacy of jules dupuit: A review essay. *History of Political Economy*, 44(3):493–504.

Eklöf, M. and Sacklén, H. (2000). The hausman-macurdy controversy: Why do the results differ across studies? *Journal of Human Resources*, pages 204–220.

Feldstein, M. (1978). The welfare cost of capital income taxation. *Journal of Political Economy*, 86(2):S29–S51.

Feldstein, M. (1995). The effects of marginal tax rates on taxable income: A panel study of the 1986 tax reform act. *Journal of Political Economy*, 103(3):551–572.

Fraker, T. and Moffitt, R. A. (1988). The effect of food stamps on labor supply: A bivariate selection model. *Journal of Public Economics*, 35(1):25–56.

Friedberg, L. (1998). The social security earnings test and labor supply of older men. *Tax policy and the economy*, 12:121–150.

Friedberg, L. (2000). The labor supply effects of the social security earnings test. *Review of Economics and Statistics*, 82(1):48–63.

Guner, N., Kaygusuz, R., and Ventura, G. (2012). Taxation and household labour supply. *The Review of Economic Studies*, 79(3):1113–1149.

Gustman, A. L. and Steinmeier, T. L. (1983). A structural retirement model. Technical report.

Gustman, A. L. and Steinmeier, T. L. (1984). Partial retirement and the analysis of retirement behavior. *ILR Review*, 37(3):403–415.

Guvenen, F. (2011). Macroeconomics with heterogeneity: A practical guide. NBER Working Paper 17622.

Hagstrom, P. A. (1996). The food stamp participation and labor supply of married couples: An empirical analysis of joint decisions. *Journal of Human Resources*, pages 383–403.

Harberger, A. C. (1964a). The measurement of waste. *American Economic Review*, 54(3):58–76.

Harberger, A. C. (1964b). Taxation, resource allocation, and welfare. In Due, J. F., editor, *The Role of Direct and Indirect Taxes in the Federal Revenue System*, pages 25–80. Princeton University Press.

Hausman, J. A. (1985). Taxes and labor supply. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of Public Economics*, volume 1, pages 213–263. Elsevier.

Hoynes, H. W. (1996). Work, welfare, and family structure: what have we learned? NBER Working Paper 5644.

Ilin, E., Kotlikoff, L., and Pitts, M. M. (2022). Does the U.S. fiscal system discourage marriage? A new look at the marriage tax. Working paper.

Iskhakov, F., Jørgensen, T. H., Rust, J., and Schjerning, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.

Kaygusuz, R. (2010). Taxes and female labor supply. *Review of Economic Dynamics*, 13(4):725–741.

Keane, M. and Moffitt, R. A. (1998). A structural model of multiple welfare program participation and labor supply. *International Economic Review*, 39(3):553–589.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles: I. the basic neoclassical model. *Journal of Monetary Economics*, 21(2-3):195–232.

Kotlikoff, L. J. (1978). Social security, time for reform. In Boskin, M. J., editor, *Federal Tax Reform: Myth or Reality*, pages 119–144. Institute for Contemporary Studies.

Lerman, S. and Manski, C. (1981). On the use of simulated frequencies to approximate choice probabilities. In Manski, C. F. and McFadden, D. L., editors, *Structural analysis of discrete data with econometric applications*, pages 305–319. MIT Press.

Lindsey, L. B. (1987). Individual taxpayer response to tax cuts: 1982–1984: With implications for the revenue maximizing tax rate. *Journal of Public Economics*, 33(2):173–206.

MaCurdy, T., Green, D., and Paarsch, H. (1990). Assessing empirical approaches for analyzing taxes and labor supply. *Journal of Human Resources*, 25(3):415–490.

Maliar, L., Maliar, S., and Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026.

Meyer, B. D. and Rosenbaum, D. T. (2001). Welfare, the earned income tax credit, and the labor supply of single mothers. *The Quarterly Journal of Economics*, 116(3):1063–1114.

Moffitt, R. A. (1983). An economic model of welfare stigma. *The American Economic Review*, 73(5):1023–1035.

Moffitt, R. A. (1992). Incentive effects of the us welfare system: A review. *Journal of Economic Literature*, 30(1):1–61.

Moffitt, R. A. (2002). Welfare programs and labor supply. In Auerbach, A. J. and Feldstein, M., editors, *Handbook of public economics*, volume 4, pages 2393–2430. Elsevier.

Moffitt, R. A. (2003). The temporary assistance for needy families program. In Moffitt, R. A., editor, *Means-tested transfer programs in the United States*, pages 291–364. University of Chicago Press.

Moore, R. and Pecoraro, B. (2020). Macroeconomic implications of modeling the internal revenue code in a heterogeneous-agent framework. *Economic Modelling*, 87:72–91.

Moore, R. and Pecoraro, B. (2021). A tale of two bases: Progressive taxation of capital and labor income. *Public Finance Review*, 49(3):335–391.

Moore, R. and Pecoraro, B. (2023). Quantitative analysis of a wealth tax for the united states: Exclusions and expenditures. *Journal of Macroeconomics*, 78:103559.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057.

Pencavel, J. (1986). Labor supply of men: a survey. In Ashenfelter, O. C. and Layard, R., editors, *Handbook of Labor Economics*, volume 1, pages 3–102. Elsevier.

Rosen, H. S. (1978). The measurement of excess burden with explicit utility funcions. *Journal of Political Economy*, 86(2):S121–S135.

Rust, J. P. (1989). A dynamic programming model of retirement behavior. In Wise, D. A., editor, *The Economics of Aging*, pages 359–404. University of Chicago Press.

Rust, J. P. (1990). Behavior of male workers at the end of the life cycle: An empirical analysis of states and controls. In *Issues in the Economics of Aging*, pages 317–382. University of Chicago Press, 1990.

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

Torczon, V. (1997). On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25.

Zabalza, A., Pissarides, C., and Barton, M. (1980). Social security and the choice between full-time work, part-time work and retirement. *Journal of Public Economics*, 14(2):245–276.