The Macroeconomic Effects of Tax Progressivity: Evidence from the Eastern European Flat-Tax Reforms[1]

Brian Wheaton[2]
UCLA, Anderson School of Management
February 25, 2023

Abstract: I study the experience of twenty post-Communist countries, which introduced flat taxation on income. I find that the flat tax reforms increase annual per-capita GDP growth by 1.33 percentage points for a transitionary period of approximately one decade. These findings are robust to multiple alternative specifications designed to deal with various identification challenges, including electoral endogeneity and plausibly-correlated reforms. Further, I find that the growth effect primarily operates through increases in investment and, to a lesser extent, labor supply. It is driven by the reductions in progressivity resulting from the reforms rather than merely reductions in the average tax rate.

1 Introduction

Between 1994 and 2011, twenty post-communist countries introduced flat taxation on incomes at varying, but typically quite low, rates. At their peak, the vast majority of Eastern European and Central Asian countries had a flat tax in effect. Since 2011, on the other hand, some of these countries have repealed their flat taxes and reverted to a progressive system of income taxation. I use these policy changes to evaluate the macroeconomic effects of tax progressivity.

Using quarterly GDP data on this panel of flat-tax adopters and a distributed-lag regression approach, I find that the adoption of a flat tax structure has a strongly significant positive effect on growth in GDP per-capita. In my preferred specification, this magnitude of this effect is approximately 1.33 percentage points annually over a decade – a cumulative effect of 14.4 pp. This result is highly robust, remaining statistically significant under a variety of alternative

[1] I am grateful to Alberto Alesina, Robert Barro, Gabe Chodorow-Reich, John Cochrane, Stefano DellaVigna, Martin Feldstein, Ed Glaeser, Jeff Miron, Gautam Rao, Stefanie Stantcheva, Matthew Weinzierl, and Paul Willen for excellent comments and advice. I am also grateful to the participants of the Harvard Macroeconomics Seminar and the Harvard PF/Labor Lunch Seminar. I kindly thank the Bradley Foundation and the Institute for Humane Studies for generous research support.
[2] Contact: brian.wheaton@anderson.ucla.edu

specifications. In particular, I control for other (potentially correlated) aspects of the business environment aside from the tax code – as measured by various components of the *Ease-of-Doing-Business Index* – to deal with the concern that the flat tax might be implemented at the same time as other reforms which are driving the effect. I drop countries which implemented a corporate tax reform or VAT reform contemporaneously with the flat income tax in order to ensure the results are indeed driven by the flat tax. I control for budget balance and lags thereof to factor out potential Keynesian stimulus effects from deficit spending that might be associated with the reforms. I restrict my analysis to the subset of flat-tax reforms passed after a close electoral victory for the party implementing the reform in order to focus on the reforms that are most plausible to be exogenous. In every single case, the effect of flat taxation on GDP growth remains statistically significant, albeit with slight changes in magnitude in some cases.

The finding of increased GDP growth is also strongly robust to a variety of alternative regression specifications, including autoregressive distributed-lag, difference-in-differences, and event-study difference-in-differences. The event study specification reveals no significant evidence of pre-trends and shows that the effect on GDP growth is transitionary and persists with statistical significance for approximately one decade. To further probe the validity of the results, I run permutation tests randomizing the timing of the flat-tax reform across countries. This alternative method of generating p-values yields even stronger statistical significance.

Next, I investigate the potential channels through which the effect on GDP growth is mediated. I find that the effect appears to be driven primarily by increased investment, with some limited (non-significant) evidence of increased employment as well. I find no evidence of increases in inequality or budget deficits – consistent with a combination of economic growth and a reduced shadow economy/increased tax compliance.

To further probe the mechanism through which the flat tax reforms have these effects, I decompose each tax reform into its constituent parts: the extent to which it reduced overall tax rates (a level shift in the tax schedule) and the extent to which it reduced progressivity (a change in the slope of the tax schedule). To do this, I utilize data on income distributions from the World Income Inequality Database coupled with tax code data from Ernst & Young and PricewaterhouseCoopers International Tax Guides. These data allow me to directly compute the change in the average marginal tax rate (AMTR) and the change in progressivity associated with each reform, fully accounting for differences in the standard deduction, personal exemptions, and differences in deductibility of payroll taxes/social contributions. I find that decreased progressivity is of greater importance for investment, employment, and GDP growth than simply reduced average marginal rates. The latter component of the reforms does not yield statistically significant effects. This attests to the importance of progressivity – rather than merely the average level of tax rates – for macroeconomic outcomes.

In order to better understand these findings, I set up a two-period model of the labor and investment decisions that individuals face under varying tax progressivity. I show that, for any concave and separable utility function, decreased progressivity has an unambiguous effect on investment and labor above and beyond the effect of a downward shift in the tax schedule alone. Intuitively, reduced marginal rates – even conditional on an individual's total tax bill – reduce the disincentive against the activity being taxed and prompt individuals to engage in more of that activity. Here, the activities in question are investment and labor. I relate these findings to a Solow model of GDP growth, which implies that a transitionary increase in growth over the short- and medium-run should result.

The magnitude of the effects on growth through these channels implied by median estimates of relevant micro elasticities is approximately one-third to one-half of the magnitude I estimate in the data. This, however, is a regularity in the literature on the effects of tax changes on GDP. As Ramey (2019) notes in a review of this literature, effect sizes implied by models calibrated with micro elasticities from the relevant literatures tend to be one-quarter to one-half of effect sizes estimated in the data using natural experiments or SVAR approaches. Furthermore, using values of the interest elasticity of saving near the top of the range of published estimates produces an effect on GDP growth consistent with (or even slightly larger than) my findings.

The structure of the remainder of this paper is as follows. In Section 2, I review some background information related to flat taxation and its macroeconomic effects – both the relevant economics literature and the political-economic context of the post-communist flat tax reforms. In Section 3, I discuss my various data sources and the empirical framework – along with its assumptions – that I use to investigate the macroeconomic effects of flat taxation. In Section 4, I extensively cover my main empirical results, a multitude of robustness checks, and an investigation of the mechanism and its consistency with the model. In Section 5, I set up and solve a two-period model of the labor and investment decisions under varying tax rates and progressivity, and I consider its macroeconomic implications through a standard Solow model of economic growth. In Section 6, I compare my findings to the effects that would be implied by micro elasticities from the existing literature. In Section 7, I conclude.

2 Political Economic Context

2.1 Post-Communist Context

For most of the latter half of the 20^{th} century, the economies of the Eastern European and Central Asian nations were centrally-planned in the Soviet design: fully state-owned and managed by bureaucratic commissions that mandated wages, prices, investment, and output through the auspices of Five-Year Plans. Following the disintegration of the Eastern Bloc and then the USSR itself in 1991, most of these countries embraced market-based reforms, which brought the economic systems of the region into greater consonance with the Western European norm. Indeed, many of these countries went even further than that in terms of reducing the role of government in economic life.

One key example of this fact is the introduction of flat income taxation. Between 1994 and 2011, twenty countries in Eastern Europe introduced such a tax, at varying – but typically quite low – rates as a percentage of income. Since 2011, on the other hand, several of these countries have repealed their flat taxes and reverted to a progressive system of income taxation. Table 1 lists the countries and the associated year of introduction/repeal of a flat income tax, along with the flat rate itself (upon introduction). Figure 1 shows these countries on a map. All of the introductions and repeals were made effective on January 1^{st} of the stated year[3].

To some extent, the blanket term "flat tax" hides the richness in variation amongst the reforms that occurred in these countries. For example, in some cases, only the income tax schedule was modified; in other cases, corporate tax, VAT, and income tax were changed simultaneously. In some cases, the tax reform was a reduction in the general level of taxation; in other cases, it was budget-balanced or even constituted an increase in the general level of taxation. Some of this information is summarized in Table 2, which shows the changes in various tax rates associated with the reform – including the average marginal rates and the standard deviation of the marginal rates (a measure of progressivity), calculated as outlined in Appendix B.

In all cases, the advocates for the reform indicated that they expected it to attract increased

[3] The sole exception is Montenegro, which made its newly-introduced flat tax effective on July 1^{st}, 2007.

investment and reduce tax evasion. Many were also influenced by the conjecture that such reforms would stimulate economic growth and thus more than pay for themselves in short order. Tax competition was another substantial motive, with Ukraine, for example, choosing its 13% rate to match that of its neighbor Russia in order to avoid being undercut and Belarus, a few years later, choosing a 12% rate in order to undercut them both. Macedonia, in 2007, chose a 10% rate in order to be the lowest in the region. The next year, its immediate neighbors Albania and Bulgaria followed at 10%. Somewhat further east, Turkmenistan introduced a 10% flat tax in 2005. Neighbors Kyrgyzstan and Kazakhstan followed at 10% in the next several years.

One potential motive of tax/expenditure changes is to use a well-defined fiscal instrument—such as a tax cut on high earners or an increase in government expenditure on infrastructure—in order to offset expected business-cycle fluctuations on the horizon. As argued by Romer and Romer (2010) and much of the subsequent literature on the macroeconomic effects of tax shocks, such endogenous tax changes are unsuitable for studying the effect of tax changes on output. While the proponents of the flat tax had strong expectations as to what its effects would be, it would have made a blunt and unlikely instrument for short-run fiscal policy. Fundamentally, the decision to introduce flat taxation in these countries was ideological in nature, often implemented after the victory of center-right coalitions, because said coalitions believed the reforms would yield desirable effects.

Furthermore, just as the implementation of such flat taxes was typically undertaken by a center-right coalition shortly after an electoral victory, their repeal typically occurred after the victory of center-left coalitions. Again, parties and individuals advocating for repeal did not make arguments based on offsetting expected forthcoming economic fluctuations. Rather, the emphasis was again ideological: concerns about fairness and disproportionate burdens on the working-class.

2.2 Literature Review

Arguably the most influential case for flat taxation was made by Hall and Rabushka (1983), who advocated for a broad-based reform to the US tax code. Their basic proposal centered on eliminating exclusions, deductions, or credits to any individual or organization and using the revenue gained to reduce marginal tax rates to a 19% flat tax on wages and business income. Notably, Hall and Rabushka propose not an income tax but a wage tax. Critics suggested that, because the Hall-Rabushka proposal was a wage tax and the vast majority of income from top earners is capital income, functionally the reform would entail a transfer from middle- and working-class individuals to wealthier individuals, with those making under $50,000 per year at 1983 prices (equivalent to approximately $140,000 per year in 2021) experiencing an increase in taxation, according to Pechman (1984). Additionally, Auerbach and Kotlikoff (1987), using tax simulation models, found that a shift from an income tax to a wage tax would actually reduce economic efficiency, suggesting that the touted efficiency gains of a flat tax on wages may not be met.

Perhaps as a result of such considerations or perhaps as a result of the anticipated political difficulties of implementing a wage tax, the Hall-Rabushka flat wage tax proposal is not precisely what has been implemented in any of the Eastern European countries. They instead feature more traditional income taxation, except at flat rates, with a standard personal deduction. As such, the above critiques do not directly apply.

There are a number of theoretical benefits of flat income taxation at low rates. First, there is a large body of evidence suggesting that there are indeed behavioral responses to income taxation, with higher rates inducing lower labor supply. The general consensus is that, for the majority of prime-age males in the United States, the earnings elasticity is rather low (in the neighborhood of 0 to 0.1), and that it has declined substantially over time for prime-age females as well (now in the neighborhood of 0.2) as they have become more attached to the labor force (Pencavel 1986, Pencavel 2002, Blau and Kahn 2007). Higher earners, however, from whom the majority of tax revenue in most systems originates, tend to be more sensitive to changes in marginal tax rates, and estimates of their earnings elasticity tend to be in the neighborhood of 0.5 to 0.8 (Saez, Slemrod, and Giertz 2012). This suggests significant benefits vis-à-vis labor supply and economic output in response to a cut in top tax rates, and it also suggests that making up some or most of the lost revenue by increasing the tax burden on low income individuals would not significantly offset those gains.

On the other hand, Rebelo and Stokey (1995) explicitly investigate the theoretical growth effects of flat taxation by calibrating an endogenous growth model to the U.S. data. They argue that flat-tax reform would have little or no effect on the U.S. growth rate. However, they note that factor shares, depreciation rates, the elasticity of intertemporal substitution, and the elasticity of labor supply are crucial parameters to which this result is sensitive, suggesting it is possible that countries with different parameter values may indeed enjoy significant economic growth effects as a result of a flat-tax reform. Simulation exercises by Fuest, Peichl, and Schaefer (2008) and Paulus and Peichl (2009) suggest positive but limited efficiency gains – usually accompanied by inequality increases – from imposing a flat tax on a Western European economy.

Another potential benefit of a flat tax with low rates – one cited frequently by flat-tax proponents in Eastern Europe – is a reduction in tax evasion, a prediction suggested by many models of the tax compliance decision. If one models evasion as a costly activity (perhaps consuming time and requiring payments to a team of "creative" accountants), then the reasoning becomes straightforward. In their micro-level study of the Russian flat-tax reform, Gorodnichenko, Martinez-Vasquez, and Peter (2009) do indeed find evidence of reduced evasion.

On the whole, though, there has been relatively little research on the macroeconomic effects of the Eastern European flat tax reforms, a surprising fact given how politically-charged the surrounding debate can be. Right-wing and left-wing commentators alike have made known their strong, even fiery, opinions on the matter[4]. However, this debate has been largely unquantified. As a result of this shortage of well-identified evidence, even a review paper on flat taxation by Keen, Kim, and Varsano (2008) is light on empirical evidence, instead focusing primarily on theoretical implications.

There are a handful of exceptions. Ivanova, Keen, and Klemm (2005) examine micro-level labor-supply responses to the Russian flat tax reform using panel data from the Russian Longitudinal Monitoring Survey, finding little to no evidence of enhanced labor-supply but substantial reductions in evasion (as measured by the gap between household expenditure and reported income). They note, however, that changes in tax enforcement accompanied the flat-tax reform, and it is difficult to decompose how much of the reduced evasion is due to this versus the flat-tax reform.

Mentioned previously, Gorodnichenko, Martinez-Vasquez, and Peter (2009) go a step further. They first supplement the Ivanova, Keen, and Klemm analysis by using the same data source but with a few extra years of data (more than the first two years after the reform), confirming the lack of any significant labor-supply/productivity response to the tax reform. However, they are able to isolate the effect of the reform on evasion from the effect of increased

[4] See, for example, Mitchell (2007), which rails against an IMF report asserting that the flat-tax reforms were unlikely to have major impacts on labor supply or tax compliance, and Bashevska (2014), which brands Macedonia a "workers' hell" and charges its flat-tax reform with increasing poverty.

enforcement by using a difference-in-difference design which takes advantage of the fact that some income brackets did not experience a marginal rate change as a result of the reform (and hence would only have experienced an enforcement change) while others did, restricting the sample to those near the marginal rate discontinuity for robustness. They find that there was indeed a strong and significant impact of the flat-tax reform on evasion, although 30% smaller than implied by the approach of Ivanova, Keen, and Klemm.

Easterbrook (2008) delves into the tax code data for eight of the countries that implemented flat tax reforms and uses said data to calculate the actual change in average marginal income tax rates. She uses these calculations to calibrate the Prescott (2004) model of labor supply for each of the countries, finding that the model predicts a substantial labor-supply increase in most of the countries—excepting two that actually experienced increased average marginal income tax rates as a result of the reform. However, when she compares the predictions of the model to actual labor supply changes (using data on hours worked from the International Labor Office), true responses appear negligible in most cases.

Adhikari and Alm (2016) use the synthetic control method to study the effect of the flat-tax reforms on the level of GDP in the case of eight specific flat-tax reforms, finding effects in each country that are positive, albeit not strongly significant. It is worth noting that these are not the same eight countries as examined by Easterbrook. Theirs is the closest existing work to this paper. However, the fact that Adhikari and Alm examine only a selection of 8 of the 25 flat-tax reforms/repeals in Eastern Europe means that any resulting pooled estimates of the GDP effects are not necessarily comprehensive/representative. Also, likely because Adhikari and Alm are more focused on short-run level effects on GDP, they do not use data on the tax code or income distribution of the countries they study, and hence they do not separate the effects of a decrease in the average marginal tax rate (shifting downward of the tax schedule) from that of flat taxation per se (flattening of the tax schedule).

World Bank (2005) examines the effects of the Slovak flat tax reform on inequality by simulating tax payments under the pre- and post-reform tax systems using household survey data. They find little to no evidence of a change in after-tax inequality resulting from the tax reform. Slovakia, however, is one of the rare flat tax countries identified by Easterbrook as having had effectively no change in average marginal tax rates as a result of the reform, rendering the result somewhat less unexpected.

Kopczuk (2012) studies the 2004 Polish tax reform. While Poland never implemented a true flat tax that applied a flat rate to all income, the 2004 reform implemented a flat tax on *business* income only. Kopczuk finds evidence of only minor reductions in revenue despite massive reductions in tax rates (from 40% to 19% for business income formerly in the highest bracket). Because the effects he finds on reported income imply elasticities substantially larger than most of the literature on taxable income elasticity, Kopczuk argues that a reduction in tax evasion is likely driving a large portion of the effects.

No paper has yet either (i) examined the effect of the flat-tax reforms on longer-term economic growth, (ii) conducted a systematic (i.e., more than single-country) examination of the mechanism of any such effect, (iii) used the full panel of flat-tax implementing and repealing countries for which data is available, or (iv) utilized these reforms to investigate the relative importance of reduced average tax rates and reduced tax progressivity for macroeconomic outcomes. I hope to fill these gaps in the literature.

3 Data and Empirical Framework

3.1 Data

To conduct my analysis, I acquire quarterly GDP data through 2016 from the Economist Intelligence Unit (EIU), which has collected said data from the Central Statistical Bureaus of the respective countries. For Albania, Kyrgyzstan, and Montenegro, the data are not available from the EIU, so I obtain it directly from the nation's Central Statistical Bureau. In some cases, the data do not come seasonally-adjusted. As such, I apply the standard x13 seasonal adjustment procedure to these series. For some of the countries in my panel, quarterly GDP data are not available before 1995, so I supplement this with interpolated annual data from the Penn World Table where said quarterly data are missing.

From the World Income Inequality Database (WIID), I obtain data on income share by decile of population and the Gini coefficient at repeated cross-sections. The WIID collates this data from numerous sources, but one source that attempts to measure such indicators in a consistent manner across almost all countries is the World Bank, so I use the World Bank estimates within the WIID. For many countries, the World Bank has annual estimates of Gini and population-by-income stretching back for decades. For other countries, the frequency is less regular. In these cases, I use the estimate from closest year.

Also from the World Bank, I obtain the World Development Indicators (WDI) dataset, which includes data on foreign direct investment, sectoral shares in the economy, population growth, and many other useful indicators. I obtain data on patents from the World Intellectual Property Organization (WIPO). From the IMF Government Financial Statistics (GFS) dataset, I obtain government budget data. Again from the Penn World Table, I obtain data on employment and annual average hours worked.

With regard to legislated changes in the tax code, I refer to and digitize information in the annual international tax guides published by Ernst & Young (*Worldwide Personal Income Tax Guide, Worldwide Corporate Income Tax Guide,* and *Worldwide VAT, GST, and Sales Tax Guide*), which detail the tax code in each country for each year since 2006. For the earlier reforms, I obtain this information from the data appendix in Easterbrook (2008), where it was collected from analogous annual tax-code reports that PricewaterhouseCoopers published at the time[5]. In a procedure described in Appendix B of this paper, I pair the tax code data with the data on income distributions in order to compute measures of the average marginal tax rate and tax progressivity (the standard deviation of the marginal tax rate). In so doing, I fully account for differences in the personal exemption, the standard deduction, and the deductibility of payroll taxes/social contributions.

Three flat-tax-implementing countries are omitted from my sample in my main regressions: Bosnia and Herzegovina, Turkmenistan, and Kyrgyzstan. Bosnia and Herzegovina is a country with a fairly unique governmental design – a consequence of the peace agreement ending the Bosnian War of the 1990s. The country is split into the Federation of Bosnia and Herzegovina, Republika Srpska, and the Brcko District. Each has a high degree of autonomy and a separate tax code. However, data on most outcomes – GDP growth, investment, etc. – are only readily available at the national level, rendering analysis difficult at best. Turkmenistan is a highly authoritarian country with severely-restricted information flows. Tax code data and quarterly GDP data do not appear to be available from any source. Finally, there is contradictory information on the year in which Kyrgyzstan's flat tax reform was implemented, with some sources claiming 2009 and others claiming 2006. In any case, the Ernst & Young tax guides do

[5] For Albania, which is missing information on its pre-reform tax code, I supplement this with IMF (2005), which provides said information. Similarly, for Macedonia before its reform, I refer to OECD (2003). I am unable to find the tax code for the year immediately prior to the reform in these cases, so I must assume that they did not change in the couple of years leading up to the reform—an assumption that holds true for the other countries in the panel.

not provide information on Kyrgystan's tax code until the 2010s, nor do the PricewaterhouseCoopers guides, preventing me from verifying the actual date[6].

3.2 Empirical Framework

For my baseline empirical specification, I utilize a modified distributed-lag (DL) approach. This entails regressing quarter-over-quarter GDP growth on multiple lags of a "flat-tax shock" variable, which equals 1 for each country in the quarter whereupon a flat tax is adopted, -1 in the quarter whereupon a flat tax is repealed, and zero otherwise. Because these are cross-country regressions, I add country and year fixed-effects to this specification to account for persistent differences in GDP growth across countries and for regional business cycles, the timing of which could potentially be correlated with adoption or repeal of a flat tax. Specifically,

$$\Delta Y_{c,t} = \alpha + \sum_{i=0}^{I} \beta_i FlatTaxShock_{c,t-i} + \gamma_c + \psi_t + \varepsilon_{c,t}\,, \qquad \text{(DL)}$$

where $\Delta Y_{c,t}$ denotes quarter-over-quarter per-capita real GDP growth in country c at time t, *FlatTaxShock* is the aforementioned flat-tax shock variable, γ_c is a country fixed-effect, ψ_t is a time fixed-effect, and $\varepsilon_{c,t}$ is the error term. Standard errors are clustered at the level of treatment: the country level.

As discussed in Section 2, the flat tax reforms tended to be adopted for ideological reasons, rather than to counteract other factors likely to influence output in the near future. This reduces concerns of systematic correlation between these tax changes and other determinants of output growth. Regardless, it could possibly be the case that individuals are more likely to vote for the center-right parties advocating flat tax reforms at certain points in their *country-specific* business cycle. Adding lags of output growth controls for the state of the economy and helps to address this possibility of policy endogeneity. Thus, I utilize a modified autoregressive distributed-lag (ARDL) approach as a robustness check:

$$\Delta Y_{c,t} = \alpha + \sum_{i=0}^{I} \beta_i FlatTaxShock_{c,t-i} + \sum_{j=1}^{J} \phi_j \Delta Y_{c,t-j} + \gamma_c + \psi_t + \varepsilon_{c,t}\,. \qquad \text{(ARDL)}$$

These mirror the main specifications used in Romer and Romer (2010) and other papers studying the macroeconomic effects of tax changes. As an alternative approach, I employ difference-in-differences specifications – one standard (DD) and one with the addition of

[6] Adding Kyrgyzstan to the main specification using either the 2009 or 2006 date does not substantially alter the result or its statistical significance.

autoregressive terms (ARDD):

$$\Delta Y_{c,t} = \alpha + \beta F_{c,t} + \gamma_c + \psi_t + \varepsilon_{c,t} \,, \qquad \text{(DD)}$$

$$\Delta Y_{c,t} = \alpha + \beta F_{c,t} + \sum\nolimits_{j=1}^{J} \phi_j \Delta Y_{c,t-j} + \gamma_c + \psi_t + \varepsilon_{c,t} \,, \qquad \text{(ARDD)}$$

where $F_{c,t}$ is a flat-tax indicator variable equal to 0 when country c does not have a flat tax system in effect at time t and equal to 1 when it does.

It is worth taking a moment to reflect upon the identification assumptions implicit in these approaches. They all rely on a parallel trends assumption: if a country which implemented a flat tax at time t had not done so, then its economic growth would have evolved along the same trajectory as those countries which actually had not yet implemented a flat tax reform by t. Because all of the countries in my sample have a similar economic history, all are from the same region, and all eventually adopt a flat tax, the assumption may be more palatable in this context. Controlling for lagged growth, as in the ARDL and DD specifications, further helps address the reverse-causality concern described above. Even after this, however, there may remain some additional concerns. Correlated policymaking is one example: what if the flat taxes tended to be passed simultaneously with other major economic reforms? To deal with these concerns and others, I conduct a number of robustness checks that account for various potential confounds by adding control variables or limiting the sample in certain ways – such as focusing exclusively on reforms implemented after close elections between a party advocating the flat tax and a party opposing it.

One way of evaluating the plausibility of the parallel-trends assumption is to check for the existence of pre-trends. Prior to adopting a flat tax, are countries on the same growth trajectory as those that will not adopt a flat tax until years later? To answer this question, I run an event-study difference-in-differences specification (sometimes called a dynamic difference-in-differences specification):

$$\Delta Y_{c,t} = \alpha + \sum\nolimits_{m=A}^{B} \beta_m F_{c,t}^{m} + \sum\nolimits_{j=1}^{J} \phi_j \Delta Y_{c,t-j} + \gamma_c + \psi_t + \varepsilon_{c,t} \,, \qquad \text{(EventDD)}$$

where $F^{m}{}_{c,t}$ denotes whether a flat tax was in it's m^{th} quarter in effect in country c during time period t. For example, the Romanian flat-tax reform was implemented in 2005Q1. Thus 2005Q2 is its second quarter in effect, 2005Q3 is its third, etc. By setting $A < 0$, this specification allows one to test for the existence of pre-trends and thereby provide evidence supporting claims about the lack of policy endogeneity, the existence of parallel trends, and the overall cleanliness of the natural experiment. This event-study specification also responds to concerns raised recently in the applied econometrics literature – such as in Borusyak and Jaravel (2017) – that running static difference-in-differences specifications over long time horizons over which treatment effects may plausibly be heterogeneous can potentially bias the resulting static regression coefficient.

These initial specifications represent all flat tax reforms with simple indicator variables. This allows readily for estimation of the average treatment effect of a flat tax reform, but it glosses over important variation across reforms. That is, due both to variation in the pre-existing tax progressivity across countries and to variation in the rate of the newly-implemented flat tax, reforms varied in the extent to which they affected both the average marginal tax rate and tax progressivity. As such, I run a revised version of my baseline DL specification with the flat-tax shock indicator variable replaced with two variables: the change in the average marginal tax rate (AMTR) and the change in the standard deviation of the marginal tax rate (SDMTR) associated with each reform/repeal:

$$\Delta Y_{c,t} = \alpha + \sum_{i=0}^{I} \mu_i \Delta AMTR_{c,t-i} + \sum_{i=0}^{I} \sigma_i \Delta SDMTR_{c,t-i} + \gamma_c + \psi_t + \varepsilon_{c,t}$$

This allows me to decompose the effect of the flat-tax reforms into (i) the effect of a downward shift in the tax schedule and (ii) the effect of a change in the tax schedule's slope.

4 Empirical Results

4.1 Main Results

The left panel of Figure 2 is a plot of the impulse response to the flat tax reforms per the distributed lag (DL) regression specification. Standard errors are clustered by country, and the plot includes 90% confidence intervals. Over a 10-year horizon, adopting a flat-tax reform is associated with a strongly-significant 14.4 percentage-point boost in GDP growth – equating to approximately 1.33pp per year. While sizeable, it is worth noting that average annual GDP growth during the 2000s in the countries in my panel was approximately 5%. Consequently, an additional 1.33pp is large but conceivable. It corresponds to a roughly 3-year boost in growth due to flat tax adoption. A different way of putting this effect size in perspective is that average annual US GDP growth was approximately 1.53pp higher during the 1990s – typically regarded as a good decade for US performance – than it was during the 2000s – typically regarded as a mediocre one.

The left panel of Figure 2 displays impulse response functions for the autoregressive distributed lag (ARDL) regression specification, with four lags of quarter-over-quarter GDP growth. Due to the autoregressive terms, cumulated standard errors for these impulse response functions are generated using Monte Carlo methods. I take 1000 draws from a multivariate normal distribution with mean and variance-covariance matrix corresponding to the point estimates and variance-covariance matrix of the regression coefficients. For each draw from the distribution, I compute the implied response of output to a flat-tax shock after q quarters, for each $q \in (0, 39)$. I then obtain the standard error for the response after q quarters by taking the standard deviation over the 1000 aforementioned implied responses after q quarters. The ARDL specification an effect size that remains in line with the baseline DL and strongly significant.

Figure 3 turns to the difference-in-differences specification. The left panel shows the results of the standard diff-in-diff (DD), which again yields a strongly significant result. The magnitude, however, is twice that of the preceding specifications. Adding the autoregressive terms (ARDD), as in the right panel, brings the magnitude back in line with the earlier distributed-lag specifications.

Figure 4 turns to the event-study difference-in-differences specification (EventDD). As pointed out by recent applied econometrics papers such as Borusyak and Jaravel (2017), coefficients from static difference-in-difference regressions may be unreliable if dynamic coefficients exhibit little stability over time. Fortunately, as can be seen in the left panel of the figure, the effect is quite stable for about a decade, and its magnitude during that time period is stable and consistent with the static specification. Pre-trends are not statistically significant and appear to bounce around zero – an encouraging sign for the parallel trends assumption. It is also worth highlighting that the fact the effect is not permanent but rather transitionary. The event-study coefficients lose statistical significance after about a decade. The right panel of the figure plots the long-term impulse response function implied by the event-study difference-in-differences coefficients. Cumulated standard errors are again obtained using Monte Carlo methods, as described above. As seen in the earlier specifications, there is sizeable and statistically-significant growth over the first decade after the reforms. From approximately year 10 onward, while the point estimate of the impulse response function continues to grow, it is not statistically distinguishable from the level attained after 10 years. In other words, there is no further statistically-significant growth induced by the flat tax reforms after the initial decade.

In addition to the assumption of parallel trends, Goodman-Bacon (2021) argues that, for causal interpretation, difference-in-differences estimates require an assumption of homogeneous treatment effects across timing of treatment (e.g., flat taxes implemented in the 1990s versus flat taxes implemented in the 2000s). Goodman-Bacon shows that any multi-period staggered-adoption difference-in-differences estimation can be expressed as a weighted average of all simple 2-by-2 (two-group, two-period) difference-in-differences in the data. This decomposition allows one to gauge the plausibility of an assumption of homogeneous treatment effects over time in the same way that an event-study approach allows one to gauge the plausibility of the parallel trends assumption.

Figure 5 shows the results of a Goodman-Bacon decomposition of the difference-in-differences (DD) specification. Note that it is necessary to drop flat-tax repeals from the sample, as Goodman-Bacon's procedure assumes the binary treatment variable (here, the flat-tax indicator variable) is weakly increasing over time. As can be seen from the figure, there is little difference between the average treatment effect amongst 2-by-2 difference-in-differences estimated by comparing early-treated units (treatment) with later-treated units (control – not yet treated) and the average treatment effect amongst 2-by-2 difference-in-differences estimated by comparing later-treated units (treatment) with early-treated units (control – already treated). In other words, treatment effects do not appear to be heterogeneous over time. Furthermore, the estimated overall treatment effect is not driven by any severe outliers. There are a few 2-by-2 difference-in-differences which constitute negative outliers, but they mostly receive low weight in the estimation. There are a couple other 2-by-2 difference-in-differences which receive higher weights than any other by a factor of two – but they are very near the average treatment effect and therefore do not drive the estimate. In summary, the results of applying Goodman-Bacon's procedure are consistent with the conditions under which a causal interpretation of the results is valid.

Table 3 summarizes the results of the DL specification and various modifications thereof, showing the cumulative effect of each after ten years and the corresponding standard errors. Column (1) repeats the baseline DL specification for purposes of comparison. Column (2) addresses the concern of correlated tax changes. The specification represented here drops from the sample countries which changed either their corporate tax code or VAT at the same time as their flat income tax was implemented; most countries in the sample do not modify their corporate tax or VAT contemporaneously with the flat income tax reform. It also includes controls for the change in payroll tax (both average marginal rate and standard deviation of the marginal rate), which I construct to include the employee's contribution rate, the employer's contribution rate, and social security contributions. Nearly all reforms were accompanied by at least a slight change in the payroll tax schedule, so it is not possible to merely drop those cases where the payroll tax was changed as one can readily do for the corporate tax and VAT. This yields very similar results, suggesting that the main result is driven not by correlated tax changes but by the flat income tax reforms themselves.

Column (3) addresses the concern of correlated policymaking beyond the realm of tax policy. The idea here is that the party introducing the flat-tax reform may also introduce correlated reforms, which could be what are actually responsible for the growth. It is worth noting that, in most all of the countries in my panel, the flat tax has been (or was) in effect for a sufficiently long time such that a different party led the government for at least as many years as the party which introduced the flat tax, somewhat reducing the magnitude of this concern. Still, in order to deal with it, I turn to the Ease-of-Doing Business Index. The Ease-of-Doing-Business Index is compiled annually by the World Bank for a panel of nearly all countries in the world. Its aim is to capture the institutional quality of the environment for starting and operating a business with 10 sub-indices[7]. Adding all of the sub-indices not pertaining to taxes to the baseline specification, I find no evidence for this conjecture, as the growth effect actually becomes slightly *larger* (though the difference is not significant) once these controls are added[8].

Column (4) restricts the sample to those countries wherein the flat tax reform was implemented after the close election victory of the party advocating flat taxation. The idea here is that a country where 80% of the populace favors a flat tax and 20% is against is plausibly a quite different place than one where support was roughly 50/50 when the issue came up for debate. Differences in outcomes amongst the latter group are thus more likely to reflect differences in policy adoption rather than idiosyncratic factors correlated with high enthusiasm for center-right policies. For the purposes of this exercise, I define a close election victory as a

[7] These are starting a business, dealing with construction permits, getting electricity, registering property, getting credit, protecting investors, paying taxes, trading across borders, enforcing contracts, and resolving insolvency.

[8] Note that the sample size is noticeably lower for this column. That is because the Ease-of-Doing-Business Index did not exist prior to 2004, limiting the sample somewhat.

victory by a margin of less than 5 percent in the most recent parliamentary election (for parliamentary republics) or presidential election (for presidential republics). The resulting sample includes Albania's flat-tax implementation (but not its repeal), the Czech Republic's implementation (but not its repeal), Latvia, Romania, Slovakia's implementation (but not its repeal), and Ukraine's repeal (but not its implementation). The fact that countries in this sub-sample are disproportionately likely to have repealed the flat tax later validates the sub-sample's character as a group of countries where support for the flat tax was much less unanimous and more actively debated. This approach yields a regression specification hinging on only six reforms – and thus likely to be under-powered. Despite this, the result remains fairly strongly significant; the magnitude of the effect actually increases substantially, though the difference is not significant.

Column (5) restricts the sample to flat tax adoptions only, dropping repeals from the sample. The magnitude of the effect is reduced slightly relative to the baseline specification, but statistical significance remains strong.

Column (6) aims to investigate whether the effects are truly an enduring consequence of the flat-tax policy itself or simply a short-term Keynesian stimulus effect that has far more to do with deficit spending from any source than the particulars of a flat tax. I add as controls to the baseline specification a measure of budget balance (deficit or surplus as a percentage of GDP) and lags of this variable. Statistical significance is retained, and the magnitude of the effect barely budges.

Column (7) adds a control for the log of per-capita GDP, in acknowledgement of the existence of convergence effects and the fact these countries tended to have less developed economies when they had progressive taxes than when they had flat taxes (since the latter is the more recent system in most of these countries). It should be noted that any bias induced by this factor should bias the effect in the main specification toward zero – downward, not upward. Regardless, the inclusion of this control scarcely changes the situation. In acknowledgement of the finding of Barro (2015) that effect sizes in regressions such as this with convergence terms may actually be biased by the inclusion of country fixed-effects, I run a version without said FEs and, in column (8), again find a significant (albeit non-significantly smaller) effect size.

A number the countries in my panel joined the European Union over my sample period. This rarely occurred at the same time as the flat tax reform, but to rule out the possibility that the effects I find are driven by accession to the E.U., I add a control for E.U. membership to the baseline specification. These results are displayed in column (9); the result is little changed.

Column (10) runs the regressions at the annual level using the Penn World Table data. To the extent that the quarterly data released by these countries are less reliable than the annual data or that the seasonal-adjustment process induces any oddities, annual data from a highly standard source should abstract from such concerns. Again, the result remains statistically-significant, although the magitude of the effect is slightly reduced relative to baseline.

Because standard errors in these regressions are clustered at the country level, the relative paucity of clusters (i.e., 17) may lead to the concern that the computed standard errors are inaccurately small and consequently over-reject the null hypothesis. To answer this concern, one can conduct randomization inference as an alternative approach to computing p-values. I conduct a placebo test whereby I randomly assign a placebo reform year to each country in the panel and randomly assign placebo repeal years to five countries in the sample. I also conduct a permutation test whereby I randomly re-assign the timing of treatment across countries in my sample. In both cases, I repeat this procedure 1000 times, plotting the implied effects of the randomization inference on the baseline DL specification in Figure 6. The left panel pertains to the placebo test, whereas the right panel pertains to the permutation test. In both cases, the implied p-value continues to suggest strong statistical significance. Indeed, here the results are actually slightly more significant than the clustered standard errors implied.

Next, I run a year-level specification analogous to the main specification, except with the change in the Gini coefficient as the dependent variable. The results are displayed in Table 4. Columns (1), (2), and (3) apply the DL, ARDL, and DD specifications, in turn. While the point estimates are small and positive – consistent with growth in inequality – they are not statistically-significant. Thus I find no evidence of any subsantial change in the level of inequality in these economies resulting from the flat tax reforms. One potential reason for this puzzling result is the fact that tax compliance was known to be very low in Eastern European and Central Asian countries prior to the flat tax reforms. If the reforms substantially boosted compliance, it would not necessarily be surprising to find a lack of any significant effect on inequality. Columns (4), (5), and (6) repeat the preceding specifications – this time with the top decile's share of total income as the outcome. The conclusion is the same.

4.2 Mechanism – Channel of the Effect

Thus flat taxation on income in Eastern Europe appears to have had a positive, robust, and rather large effect on economic growth. A key question remains: through what economic channel(s) was this effect realized? Theory and the assertions of Eastern European flat-tax proponents suggest a few possibilities:

- Domestic Investment: Reduction of tax rates and progressivity may motivate individuals to engage in activities that are likely to boost future income, including re-allocating more of their current income toward saving/investment.
- Labor Supply: Reduction of tax rates and progressivity may motivate individuals (particularly high-income individuals) to supply more of their own labor and thus generate more economic output. Such a level effect is theoretically straightforward and well-founded. An effect on economic growth rates through this channel could additionally be realized through an endogenous growth framework.
- Foreign Direct Investment: Some Eastern European proponents of flat-taxation suggested it would attract foreign investors to their countries, persuading said individuals to invest, start a business, and move there, bringing themselves along with their financial interests. Such investment could spur economic growth.
- Systematic Budget Deficit: Most of the reforms represented a mechanical reduction in the general level of taxation. If government expenditure was not reined in by a commensurate amount or if compliance was not boosted, it could be the case that the flat-tax reforms have represented systematic budget deficits, which – viewed as repeated Keynesian stimuli – could result in debt-fueled (and therefore likely unsustainable) economic growth.
- Shadow Economy Size: A notable characteristic of the Eastern European economies is the extremely large size of their underground/shadow sectors, estimates of which tend to be in the range of 40-50%, depending on the country. If reducing marginal tax rates on high-income individuals in these countries made it cheaper and easier to simply report one's income and pay one's taxes than to hire a team of "creative accountants", then measured economic growth may at least partially reflect movement of the shadow sector out of the shadows.
- Innovation: Reduction of tax rates and progressivity may cause individuals to dedicate

more effort to innovation and technical progress if they believe a higher fraction of the gains from this progress will accrue to them personally.

- Removal of Sectoral Distortions: A key feature of the Communist-era Eastern European economies was an inordinately high share of heavy industry in the overall economy. Furthermore, member states of the CMEA – the Communist equivalent of the European Economic Community – were strongly encouraged to specialize in certain areas (e.g., Romania was directed to specialize in agriculture, East Germany in tech, etc.). If an environment of high taxes/subsidies in the aftermath of this period kept sectoral allocation distorted, transition to a low, flat-tax regime could plausibly induce economic growth by remedying this situation.

All of the aforementioned hypotheses have testable implications and can be addressed here. In Table 5, I run difference-in-differences regressions precisely analogous to the main specification, albeit with differing left-hand-side variables. First, column (1) examines the effect on investment growth. I find evidence of a large and strongly statistically-significant increase of 29% over a ten-year horizon, translating to approximately 2.6% per year.

Turning next to labor supply, the flat-tax reforms could potentially have had an effect on the extensive margin or the intensive margin. As can be seen in columns (2) and (3) of Table 5, there is some weak evidence of an increase in both the employment-to-population ratio and average hours worked – a cumulative ten-year increase of 1.82pp in the former (which equates to 3.6% of the mean employment-to-population ratio) and 47 extra hours of work in a year in the latter (which equates to 2.5% of the mean average hours worked). These effects, however, are not significant at conventional levels.

Column (4) examines foreign direct investment (FDI). No remotely significant effect is found here. Although a potential effect on FDI was touted by some Eastern European flat-tax advocates, such an effect would have to occur through a much more circuitous pathway. For example, US citizens who invest money in an Eastern European country would still need to pay some US taxes on any income resulting from such investments, unless they became a resident of the country in which they are investing – a very hefty and costly decision.

Column (5) turns to the budget balance, measured as a percentage of GDP (positive in the case of a surplus, negative in the case of a deficit). The hypothesis of the GDP growth being driven by repeated Keynesian stimuli would necessitate a negative effect of the reform on budget balance. On the contrary, evidence is actually found of a positive effect. Budgetary concerns were cited by some Eastern European flat-tax opponents, but these do not appear to have been borne out. The increased output resulting from the reforms likely ameliorated direct revenue decreases.

Column (6) turns to the matter of the shadow economy, which is a sizeable component of economic activity in many Eastern European and Central Asian countries. I use estimates of shadow economy size from Hassan and Schneider (2016). Schneider has produced the most well-recognized, well-cited estimates of shadow economy size in the literature, and the 2016 update of this dataset covers the period 1999 – 2013 for nearly all countries, which overlaps with the adoption (and repeal) of the vast majority of flat taxes in my panel. These estimates are imperfect, but for countries where more accurate estimates based on the tax gap can be calculated, they match very closely with the Schneider data. In column (6), shadow economy share – the fraction of economic activity esimated to be due to the shadow sector – is used as the outcome variable. While the point estimate I find suggests a reduction in the size of the shadow sector – which would accord with conjectures that the reforms improved tax compliance – this result is not statistically significant.

Column (7) analyzes World Intellectual Property Organization (WIPO) annual data on patent applications by country. While the point estimate suggests a slight boost in patenting activity is associated with the flat tax reforms, the result is not remotely statistically-significant, and thus it cannot be said that the flat tax reforms are leading to an explosion of innovation, at least as measured by patent data[9].

Column (8) examines the sectoral distortion hypothesis, its implication is that the introduction of flat taxation would result in systematically higher structural change. The canonical method for measuring structural chance is to use the Lilien Index, named for Lilien (1982), which measures structural change by summing squared changes in the employment share of each sector, weighted by that sector's size as a fraction of total employment. Applying this technique to three-sector (agriculture, industry, services) data on employment shares in column (8), no statistically-significant effect of flat taxation on structural change is found. It is worth noting that if the structural change is occurring at a finer level (e.g., workers in the chemical

[9] An alternative specification which analyzes patents which were granted, not merely patents which were filed, similarly yields a non-significant positive coefficient.

industry becoming workers in the metal industry), it would not be detected by these measures. Regardless, the key distortion of the Communist-era economies was excessive industry and insufficient services, so one might expect movement along that margin, which would indeed be picked up by these measures.

4.3 Mechanism – Tax Level, Tax Progressivity, or Both?

Is the increased economic growth a result of the fact that the flat-tax reforms reduced average rates or of the fact that they reduced progressivity, flattening the whole tax schedule? To investigate this question, I use the *Ernst & Young* data on annual tax schedules and the WIID data on income distributions in a procedure described in Appendix B to compute the average marginal tax rate (a measure of the average level of the tax schedule) and the standard deviation of the marginal tax rate (a measure of the progressivity of the tax schedule) before and after each flat-tax reform in order to use variation in these variables stemming from the reforms. The average marginal tax rate (AMTR) is quite standard and has a long history in the literature on taxation, dating back to Barro and Sahasakul (1983, 1986). The standard deviation of the marginal tax rate (SDMTR) is a natural extension which measures progressivity. For example, a country with a standard deviation of the marginal tax rate equal to zero is a country with a flat tax. The higher the value of this standard deviation, the more the marginal tax rate varies across individuals – i.e., the more progressive the tax schedule[10].

I run a revised version of my baseline specification with the flat-tax shock indicator variable replaced with two variables: the change in the average marginal tax rate (AMTR) and the change in the standard deviation of the marginal tax rate (SDMTR) associated with each reform/repeal. This allows me to decompose the effect of the flat-tax reforms into (i) the effect of a downward shift in the tax schedule and (ii) the effect of a change in the tax schedule's slope. As in the baseline specification, I include lags of GDP growth, country fixed-effects, and year-quarter fixed-effects. The results are given in column (1) of Table 6. It can be seen that decreasing SDMTR increases GDP growth. A decrease in AMTR appears unimportant for growth; the coefficient on AMTR is close to zero and not statistically significant. The impulse response

[10] In theory, a non-zero standard deviation of the marginal tax rate could represent either a progressive tax schedule wherein low-income individuals pay a lower tax rate than high-income individuals *or* regressive tax schedule wherein low-income individuals pay a higher tax rate than high-income individuals. In practice, tax rates in my sample are (weakly) monotonically increasing in income. As such, a higher value of the standard deviation of the marginal tax rate represents a higher level of progressivity.

function corresponding to this specification is displayed in Figure 7. Columns (2), (3), and (4) repeat this regression for the investment growth, extensive-margin labor supply growth, and labor supply growth outcomes. Here, too, it is SDMTR – the progressivity of the tax schedule – that matters.

5 Model: Labor and Investment with Varying Tax Progressivity

To help understand these effects, I present a model of the labor/leisure and consumption/investment decision under varying tax progressivity. Individuals live for two periods. They obtain utility from consumption and disutility from labor. In the first period, they choose how much of their own labor, L, to supply at wage, w. They then choose how much of their resulting income to consume in the first period and how much to invest in a (risk-free) asset with return R which allows them to transfer their wealth to the next period. At the start of the second period, they pay taxes on both their labor and interest income, and they consume their remaining wealth. The tax rate is given by

$$\tau(y) = \alpha + \beta y$$

for any level of total income, y. Notice that, for any $\beta > 0$, the income tax rate is increasing in income – the standard definition of a progressive tax schedule. When $\beta = 0$, the income tax rate reduces to the flat base rate of α (the standard setup in models of investment under taxation). It can be shown that, with only a very weak and standard assumption on the form of the utility function, decreased tax progressivity unambiguously decreases investment and labor supplied.

Proposition: Consider a two-period model of labor and investment. Individuals choose how much labor, L, to supply at wage, w, in the first period. They then choose how much of their resulting income to consume, c_1, in the first period and how much to invest in an asset with return r for consumption in the next period. At the beginning of the second period, total income, $y \equiv wL + r_*(wL - c_1)$, is taxed at the rate $\tau(y) = \alpha + \beta y$, where $\alpha, \beta > 0$, and the marginal rate is never above 100%. In this framework, for any separable utility function u satisfying $u_c > 0$, $u_{cc} < 0$ and $u_L < 0$, $u_{LL} < 0$, the individual's chosen labor supply, L, and level of investment, $S \equiv wL - c_1$,

(i) decline, if risk aversion is sufficiently low, as the flat base rate α increases (holding constant progressivity β) AND

(ii) decline *unambiguously* as the progressivity β of the tax schedule increases (holding constant the average tax rate τ).

This proposition is proven in Appendix A of the paper. In short and intuitive terms, it says that while shifting the whole tax schedule downward may only lead to boosted investment and labor supply under certain circumstances, changing its slope (even conditional on the average tax rate) *unambiguously* boosts investment. This result is important because it indicates that progressivity of the tax schedule has implications above and beyond simple changes in the base average marginal tax rate. If changes in the income tax schedule have any effect on investment, this result suggests that flat-tax reforms are precisely where we would be most likely to detect them.

But what are the implications for economic growth, if any? To answer this question, I consider a simple Solow (1956) growth model. In the Solow model, net investment in a period is equal to output in that period times the saving rate minus any depreciation of the pre-existing capital stock. That is, $\dot{K}_t = sY_t - \delta K_t$, where Y_t denotes aggregate output in period t, K_t denotes the capital stock, s denotes the saving rate, and δ denotes depreciation. In the steady-state of the model net investment is zero, as $sY_t = \delta K_t$. Increased savings on the part of the populace lead to positive net investment and movement to a higher steady-state level of the capital stock. Since output in the Solow model is produced according to the aggregate production function $Y_t = K_t^{\alpha}(AL_t)^{1-\alpha}$, this translates to a higher level output. This higher level of output, however, is not realized instantaneously. In each period, the new, higher level of investment expands the capital stock slightly, expanding productive capacity and thus output along with it. As the capital stock continues to expand, total depreciation in each period also increases. Eventually, the higher level of depreciation catches up with the higher level of investment and the economy settles at its new steady-state. In this manner, there is a *transitionary*, non-permanent increase in economic growth as a result of the tax reform. Figure 8 plots a graphical example of this dynamic.

It should be noted that, in an intuitive sense, this framework is relatively broad in its applicability. That is, while literal financial investment in an asset which pays some return is the most obvious form of investment, a broad class of actions fit under the umbrella of sacrificing utility in the present period in order to obtain an improved payoff in the latter period. To the extent that working harder at a job increases one's future income (through promotions) or to the extent that getting a higher education improves one's future income, the above framework is applicable with only mild adjustment. These, too, are "investments", and while financial investment may be the foremost amongst them, the takeaway is that a broad set of economic activities may be affected by tax progressivity. And to the extent this broad class of additional variables affect economic growth – through endogenous growth models in the case of increased labor or human capital models in the case of improved education – limiting one's viewing lens to standard investment through the Solow model may yield but a lower bound on the importance of tax progressivity.

6 Effect Size and Elasticities

How do the effect sizes of the tax reforms that I uncover in the data compare to effect sizes implied by estimates of micro elasticities in the literatures on responsiveness of labor to changes in the after-tax wage and responsiveness of saving/private investment to changes in the after-tax interest rate? Chetty et al. (2013) conduct a meta-analysis of the literature estimating the aggregate-hours Hicksian elasticity of labor supply, finding a value of 0.50 based on macro evidence or 0.58 based on micro evidence. The mean decrease in the average marginal tax rate associated with the flat tax reforms is 5.97 percentage points. This implies an increase in labor supply of approximately 2.99 to 3.46%.

Elmendorf (1997) reviews the literature estimating the interest elasticity of saving – an area of research that was particularly active in the 1980s and 1990s. He notes that estimates in the literature are very widely-dispersed, but suggests 0.50 as a preferred estimate resulting from a standard lifecycle model "with empirically-supported parameters." Saving is disproportionately undertaken by the highest-income indviduals in society. Estimates by Dynan, Skinner, and Zeldes (2004) imply that the majority of saving in the US economy is undertaken by the top 5% of earners. In high-inequality transition economies, this distribution is likely to be even more unequal. The mean decrease in the top marginal tax rate associated with the flat tax reforms is 13.61 percentage points[11]. Consequently, using the 0.5 elasticity, an optimistic estimate of the increase in saving would be 6.81%.

What would an increase in labor supply of 3.46% and an increase in saving of 6.81% imply for GDP growth? Across countries and time, most published estimates of capital elasticity range

[11] The mean is not driven by outliers; the median is 14 percentage points.

from 0.3 to 0.5[12]. Diminishing marginal returns suggest that transition economies which have recently undergone massive losses of capital stock are likely to be at the high end of this range. So, if production is constant returns-to-scale and if the capital elasticity of output is 0.5, an increase in labor supply of 3.46% and an increase in saving of 6.81% would translate into approximately 5.1% total GDP growth – between one-third and one-half the size of most of my estimates of the cumulative effect of the flat tax reforms found in Table 3.

Thus, there is a discrepancy between the effect on GDP growth that I find in the data and the effect implied by coupling micro elasticities from the literature with the size of the flat tax reforms. As discussed by Ramey (2019), however, this discrepancy is something of a regularity in the literature. Across countries and time, the effect of tax changes on GDP growth implied by models calibrated using micro elasticities is consistently one-half or less the size of the effect estimated in the data using natural experiments or SVAR approaches.

Furthermore, it is worth noting that some work on the interest elasticity of saving has found elasticities several times higher than the 0.5 value used above. Summers (1981) and Evans (1983), for example, find values of 2.38 and 2.97, respectively. Using the 2.38 elasticity from Summers would yield a 32.4% increase in investment – slightly higher than the number I find in the data. Again using constant returns-to-scale production and a capital elasticity of 0.5, this would then translate into a total effect on GDP growth of approximately 17%, a figure only slightly larger than the effect found in most of my regression specifications. In short, it is possible to rationalize the effect sizes I uncover, but they imply large values of certain parameters – primarily the interest elasticity of saving.

7 Conclusion

Between 1994 and 2011, the spectre of flat-taxation haunted Eastern Europe and Central Asia – and, despite flat-tax repeals in several countries, flat income taxation remains in effect in most of the countries that introduced it during that era. The results of the analysis here demonstrate that flat income taxation had significant, robust, and economically large effects on GDP growth – an annualized 1.33 percentage-point effect in the main specification, which is a modified distributed-lag specification controlling for country and year fixed-effects. Although

[12] See, for example, Aschauer (1989), Levy (1990), and Berndt and Hansson (1992). Boskin and Lau (1990) find a lower capital elasticity of around 0.2-0.25. Romer (1987) finds a higher capital elasticity of 0.7-1.0.

the effect varies somewhat depending on the precise specification used, it is always strongly significant, and it is found to endure for approximately one decade. Robustness checks aimed at controlling for the possibility that parties which introduce flat taxes are conceivably more likely to foster a pro-growth environment in other ways, controlling for electoral endogeneity with a restriction of the panel to countries where the flat-tax was introduced (or repealed) after a close electoral victory, and combating potential econometric bias all retain strong significance of the aforementioned effect. Finally, deeper analysis of the channels through which the growth rate effect could possibly proceed reveals that domestic investment is the key element, with labor supply and potentially movement of economic activity out of the shadow sector playing secondary roles. However, no evidence is found for increased FDI, systematic budget deficit, or removal of sectoral distortions as a result of the flat-tax reforms.

Decomposing the flat-tax reforms into a reduction in the average marginal tax rate and a reduction in progressivity (the standard deviation of the marginal tax rate), I find that the role of progressivity is more important. In other words, in terms of boosting investment and (transitionary) economic growth, tax progressivity matters above and beyond simply the average level of the tax rate.

To better understand these findings, I set up a two-period model of the household labor and investment decisions under varying tax progressivity. I show that, for any concave and separable utility function, decreased progressivity has an unambiguous effect on investment and labor above and beyond the effect of a downward shift in the tax schedule alone. Relating these findings to a Solow model of GDP growth implies that a transitionary increase in growth over the short- and medium-run should result. Coupling investment and labor elasticities from the micro literature with the actual tax rate changes associated with the flat tax reforms suggests an effect size on GDP growth approximately one-half to one-third of the magnitudes I find – a discrepancy, but a discrepancy that is consistent with existing work on the effects of tax changes on GDP.

The extent to which these findings have applicability outside of Eastern Europe is certainly open to discussion. On the one hand, all of these countries had very similar shared histories over the course of the past three-quarters of a century: being devastated by World War II, then transformed into a Communist-led planned economy, and finally beginning a turmoil-ridden transition to market economics in the early 1990s. Because developed Western economies did not suffer from massive amounts of capital depreciation in the 1990s, it is possible they may not have quite as much to gain from boosts to capital accumulation. On the other hand, one could argue that the developing world does indeed have much to gain from such a boost. As such, a potential avenue for fruitful future research could be examining the effects of flat taxation in the developing countries of Latin America and Africa where such taxes have recently begun to be adopted.

References

Adhikari, B. and J. Alm (2016). "Evaluating the Economic Effects of Flat Tax Reforms Using Synthetic Control Methods." *Southern Economic Journal* 83(2): 437-463.

Aschauer, D. A. (1989). "Is Public Expenditure Productive?" *Journal of Monetary Economics*, 23(2): 177-200.

Auerbach, A.J. and L. Kotlikoff (1987). *Dynamic Fiscal Policy*. Vol. 11. Cambridge: Cambridge University Press.

Bashevska, M. (2014). "Macedonia: Tax Haven – and Workers' Hell." *LeftEast*. LeftEast, 18 Sep 2014. Web. 20 May 2017.

Barro, R.J. (2015). "Convergence and Modernisation." *The Economic Journal*, 125(585): 911-942.

Barro. R.J. and C. Sahasakul (1983). "Measuring the Average Marginal Tax Rate from the Individual Income Tax." *Journal of Business* 56: 419-457.

Barro. R.J. and C. Sahasakul (1986). "Average Marginal Tax Rates from Social Security and the Individual Income Tax." *Journal of Business* 59: 555-566.

Borusyak, K. and X. Jaravel (2017). "Revisiting Event Study Designs." Working Paper.

Berndt, E.R. and B. Hansson (1992). "Measuring the Contribution of Public Infrastructure Capital in Sweden." *Scandinavian Journal of Economics* 94: 151-168.

Blau, F. and L. Kahn (2007). "Changes in the Labor Supply Behavior of Married Women: 1980-2000." *Journal of Labor Economics*, 25: 393-438.

Boskin, M.J. and L.J. Lau (1990). "Post-War Economic Growth in the Group-of-Five Countries: A New Analysis." NBER Working Paper No. 3521.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2013). "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities." *NBER Macroeconomics Annual*, 27(1): 1-56.

Clementi, F. and M. Gallegati (2005). "Pareto's law of income distribution: Evidence for Germany, the United Kingdom, and the United States." *Econophysics of Wealth Distributions*: 3-14. Milan: Springer Milan.

Dynan, Skinner, and Zeldes (2004). "Do the Rich Save More?" *Journal of Political Economy* 112(2): 397-444.

Easterbrook, K.F. (2008). “Flat Taxes and Labor Supply in Central and Eastern Europe.” Senior Honors Thesis, Stanford University.

Elmendorf, D. W. (1996) “The Effect of Interest-Rate Changes on Household Saving and Consumption: A Survey.” Federal Reserve Board, Finance and Economics Discussion Series Working Paper no. 1996-27.

Evans, O. J. (1983). “Tax Policy, the Interest Elasticity of Saving, and Capital Accumulation: Numerical Analysis of Theoretical Models.” *American Economic Review*, 73(3): 398-410.

Fuest, C., Peichl, A., and Schaefer, T. (2008). “Is a Flat Tax Reform Feasible in a Grown-up Democracy of Western Europe? A Simulation Study for Germany.” *International Tax and Public Finance*, 15(5): 620-636.

Goodman-Bacon, A. (2021). “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics*, 225(2): 254-277.

Gorodnichenko, Y., J. Martinez-Vazquez, and K.S. Peter (2009). “Myth and Reality of Flat Tax Reform: Micro Estimates of Tax Evasion and Productivity Response in Russia.” *Journal of Political Economy* 117: 504-554.

Hassan, M. and F. Schneider (2016). “Size and Development of the Shadow Economies of 157 Countries Worldwide: Updated and New Measures from 1999 to 2013.” IZA Discussion Paper No. 10281.

Hall, R.E. and A. Rabushka (1983). *Low Tax, Simple Tax, Flat Tax*. New York: McGraw-Hill Companies.

IMF (2005). “Albania: Selected Issues and Statistical Appendix.” IMF Country Report No. 05/90.

Ivanova, A., M. Keen, and A. Klemm (2005). “The Russian Flat Tax Reform.” *Economic Policy*, 20: 397-444.

Keen, M., Y. Kim, and R. Varsano (2008). “The ‘Flat Tax(es)’: Principles and Experience.” *International Tax and Public Finance*, 15: 712-751.

Kopczuk, W. (2012). “The Polish Business ‘Flat’ Tax and Its Effect on Reported Incomes: A Pareto Improving Tax Reform.” Columbia University Working Paper.

Levy, D. (1990). “Aggregate Output, Capital, and Labor in the Post-war U.S. Economy.” *Economics Letters* 33(1): 41-45.

Lilien, D. M. (1982). “Sectoral Shifts and Cyclical Unemployment.” *Journal of Political Economy*, 90 (4): 777–793.

Mitchell, D.J. (2007). “The IMF’s Remarkably Shoddy Flat Tax Study.” *FoxNews.com*. Fox

Entertainment Group, 5 Jan 2007. Web. 20 May 2017.

OECD (2003). "Tax Policy Assessment and Design in Support of Direct Investment–A Study of Countries in South East Europe." *Stability Pact: SEE Compact for Reform, Investment, Integrity, and Growth*. OECD Directorate for Financial, Fiscal, and Enterprise Affairs.

Paulus, A. and Peichl, A. (2009). "Effects of Flat Tax Reforms in Western Europe." *Journal of Policy Modeling*, 31(5): 620-636.

Pechman, J.A. (1984). "Review: 'Low Tax, Simple Tax, Flat Tax' by Robert E. Hall; Alvin Rabushka." *Journal of Political Economy,* 92(2): 340-342.

Pencavel, J. (1986). "Labor Supply of Men: A Survey." *Handbook of Labor Economics,* Vol. 1.

Pencavel, J. (2002). "A Cohort Analysis of the Association between Work Hours and Wages among Men." *The Journal of Human Resources,* 37(2): 251-274.

Prescott, E.C. (2004). "Why Do Americans Work So Much More than Europeans?" *Federal Reserve Bank of Minneapolis Quarterly Review,* 28(1): 2-13.

Romer, P. M. (1987). "Crazy Explanations for the Productivity Slowdown." *NBER Macroeconomics Annual*, 2: 163-202.

Romer, C.D. and D.H. Romer (2010). "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Schocks." *American Economic Review*, 100(3): 763-801.

Saez, E., J. Slemrod, and S.H. Giertz (2012). "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature,* 50(1): 3-50.

Sinelnikov-Murylev, S.G. et al. (2003). "Analysis of Certain Results of the Personal Income Taxation Reform." *Russian Economy in 2002: Trends and Perspectives*, 24.

Summers, L. H. (1981). "Capital Taxation and Accumulation in a Life Cycle Growth Model." *American Economic Review* 71(4): 533-544.

Ramey, V. A. (2019). "Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?" *Journal of Economic Perspectives*, 33(2): 89-114.

Rebelo, S. and N.L. Stokey (1995). "Growth Effects of Flat-Tax Rates." *Journal of Political Economy*, 103(3): 519-550.

World Bank (2005). "The Quest for Equitable Growth in the Slovak Republic." World Bank Report No. 32433_SK. Washington: World Bank

Tables and Figures

Table 1: Flat Tax Timing and Rates

Country	Year of Introduction (Repeal)	Flat-Tax Rate
Estonia	1994	26%
Lithuania	1994	33%
Latvia	1997	25%
Russia	2001	13%
Serbia	2003	12%
Slovakia	2004 (2013)	19%
Ukraine	2004 (2011)	13%
Georgia	2005	20%
Romania	2005	16%
Turkmenistan	2005	10%
Kazakhstan	2007	10%
Macedonia	2007	10%
Montenegro	2007 (2013)	15%
Albania	2008 (2014)	10%
Bulgaria	2008	10%
Czech Republic	2008 (2013)	15%
Belarus	2009	12%
Bosnia & Herzegovina	2009	10%
Kyrgyzstan	2009	10%
Hungary	2011	15%

Table 2: Tax Rates and Reform Changes

Variable	Average Level	Average Δ with Reform
Income AMTR	18.2	-4.0
Income S.D. MTR	2.7	-2.2
Payroll AMTR	36.8	-1.1
Payroll S.D. MTR	2.6	+0.4
Corporate Tax Rate	20.6	-2.4
VAT Rate	19.1	+0.18

Note: Payroll includes employees' component and employers' component of payroll taxes along with social security contributions. AMTR and SDMTR are as calculated in Appendix B.

Table 3: Main Specifications

Dependent Variable: Per-Capita GDP Growth	(1) Baseline	(2) Other Tax Reforms Omitted	(3) *Doing Business*	(4) Close Elections	(5) Repealers Omitted	(6) Fiscal Size Controls	(7) Converg-ence	(8) No Country FEs	(9) E.U. Member Control	(10) Annual
Cumulative 10-Yr Effect	14.43***	14.25***	14.89**	25.97**	13.64***	15.99**	14.23***	13.35***	14.65***	12.96***
	(3.510)	(4.446)	(5.817)	(10.90)	(4.126)	(6.817)	(3.681)	(3.393)	(3.559)	(3.799)
Observation Frequency	Quarterly	Quarterly	Quarterly	Quarterly	Quarterly	Quarterly	Quarterly	Quarterly	Quarterly	Annual
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
Time FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	969	513	608	305	684	644	969	969	969	255

Note: * $0.05 < p \leq 0.1$, ** $0.01 < p \leq 0.05$, *** $p \leq 0.01$. Column (1) studies the effects of a flat tax reform on per-capita GDP growth using the DL specification given by Equation (1). Result suggests flat tax reforms are associated with a cumulative 14.57% boost in per-capita GDP over a decade, corresponding to roughly a 1.38% annual increase in growth. Column (2) omits countries with a simultaneous corporate tax reform or VAT reform from the sample and controls for the change in payroll AMTR and payroll SDMTR (these cannot merely be dropped from the sample, as nearly all reforms modify the payroll tax schedule slightly). Column (3) controls for various sub-indices of the Ease-of-Doing-Business index to factor out other reforms plausibly correlated with tax reform. Column (4) restricts the sample to countries where there was a close election prior to the reform. Column (5) controls for budget balance (deficit or surplus as percent of GDP). Column (6) controls for the log of per-capita GDP to account for convergence. Column (7) drops country fixed-effects from the Column (6) specification to deal with potential Hurwicz-Nickell bias as suggested by Barro (2015). Column (8) represents the baseline specification with annual data from the Penn World Table. Column (9) adds the log of per-capita GDP as a control and drops country fixed-effects.

Table 4: Inequality Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	Gini (DL)	Gini (ARDL)	Gini (DD)	Top Decile Share (DL)	Top Decile Share (ARDL)	Top Decile Share (DD)
Cumulative 10-Yr Effect	0.865	0.869	0.957	0.313	0.354	0.675
	(1.286)	(1.502)	(3.108)	(0.998)	(1.205)	(2.510)
Observation Frequency	Annual	Annual	Annual	Annual	Annual	Annual
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	255	255	372	255	255	372

Note: * $0.05 < p \leq 0.1$, ** $0.01 < p \leq 0.05$, *** $p \leq 0.01$. Column (1) studies the effect of a flat tax reform on inequality as measured by the Gini coefficient using the distributed-lag specification. Column (2) represents the autoregressive distributed-lag specification. Column (3) represents the difference-in-differences specification. Columns (4), (5), and (6) repeat the preceding three with a different outcome: the top decile's percentage share (out of 100) of income.

Table 5: Channel-of-Effect Specifications

Dependent Variable:	(1) Investment Growth	(2) Δ EPOP Ratio	(3) Δ Hours Worked	(4) FDI Growth	(5) Budget Balance	(6) Δ Share Shadow	(7) Patenting Growth	(8) Lilien Index
Cumulative 10-Yr Effect	29.38**	1.818	46.65	54.48	5.080*	-1.996	0.578	-0.252
	(12.34)	(1.322)	(85.82)	(53.83)	(2.742)	(1.539)	(0.586)	(0.323)
Observation Frequency	Annual	Annual	Annual	Quarterly	Annual	Annual	Annual	Annual
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	255	289	135	455	272	266	270	289

Note: * $0.05 < p \leq 0.1$, ** $0.01 < p \leq 0.05$, *** $p \leq 0.01$. Each column in Table 6 corresponds to the baseline DL specification from Table 3, albeit with a different left-hand-side variable. Column (1) studies the percent growth in investment. Column (2) studies the change in the employment-to-population ratio. Column (3) studies the change in hours worked. Column (4) studies the percent growth in foreign direct investment (FDI). Column (5) studies budget balance – deficit or surplus as a percent of GDP. Column (6) studies the change in the share of the economy estimated to be part of the shadow sector by Hassan and Schneider (2016). Column (7) studies the percent growth in patent applications. Column (8) studies the Lilien (1982) index of structural change.

Table 6: AMTR & SDMTR Regressions

Dependent Variable:	(1) GDP Growth	(2) Investment Growth	(3) Δ EPOP Ratio	(3) Δ Hours Worked
Cumulative 10-Yr AMTR Effect	-0.140	-0.893	0.016	3.373
	(.436)	(1.327)	(.136)	(5.971)
Cumulative 10-Yr SDMTR Effect	-1.923**	-6.879**	-0.623**	-26.695*
	(0.759)	(2.829)	(0.239)	(14.386)
Observation Frequency	Quarterly	Annual	Annual	Annual
Country FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	969	255	289	135

Note: * $0.05 < p \leq 0.1$, ** $0.01 < p \leq 0.05$, *** $p \leq 0.01$. Column (1) corresponds to the baseline specification from Table 3, albeit with the flat-tax shock indicator variable replaced with two variables: the change in the average marginal tax rate (AMTR) and the change in the standard deviation of the marginal tax rate (SDMTR) – a measure of the progressivity of the tax code – associated with each reform/repeal. Columns (2) - (4) repeat this exercise with different outcome variables.

Figure 1: Flat Tax Reform Map

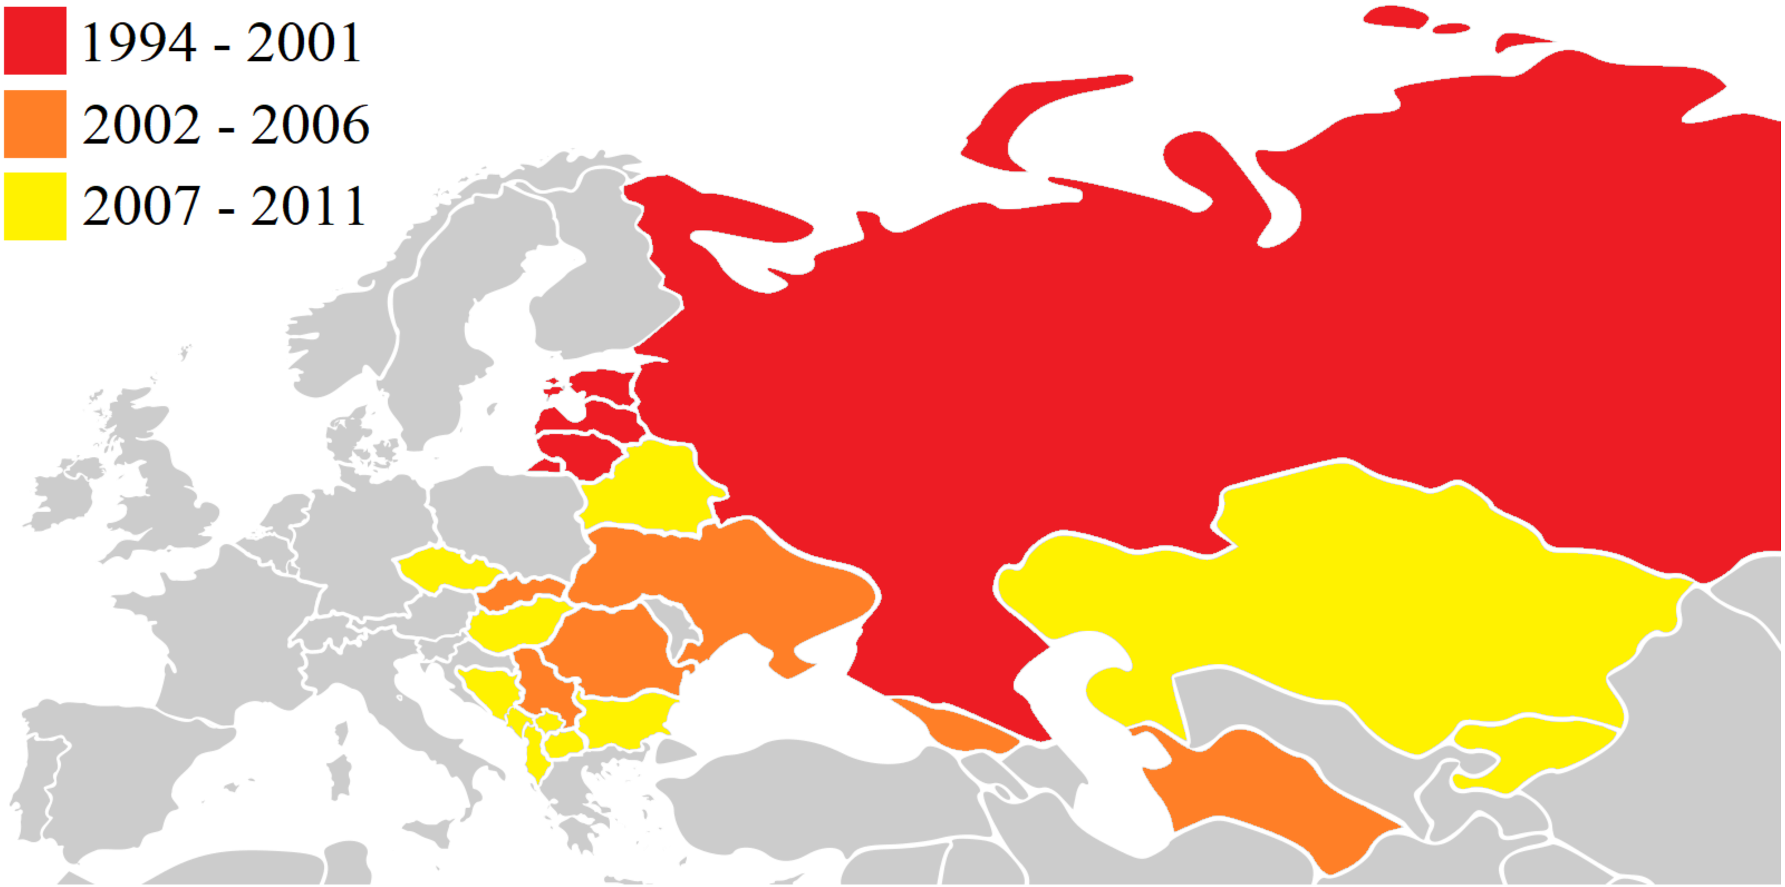

Figure 2: Effect of Flat Tax Reform on GDP Growth – Distributed-Lag Specifications

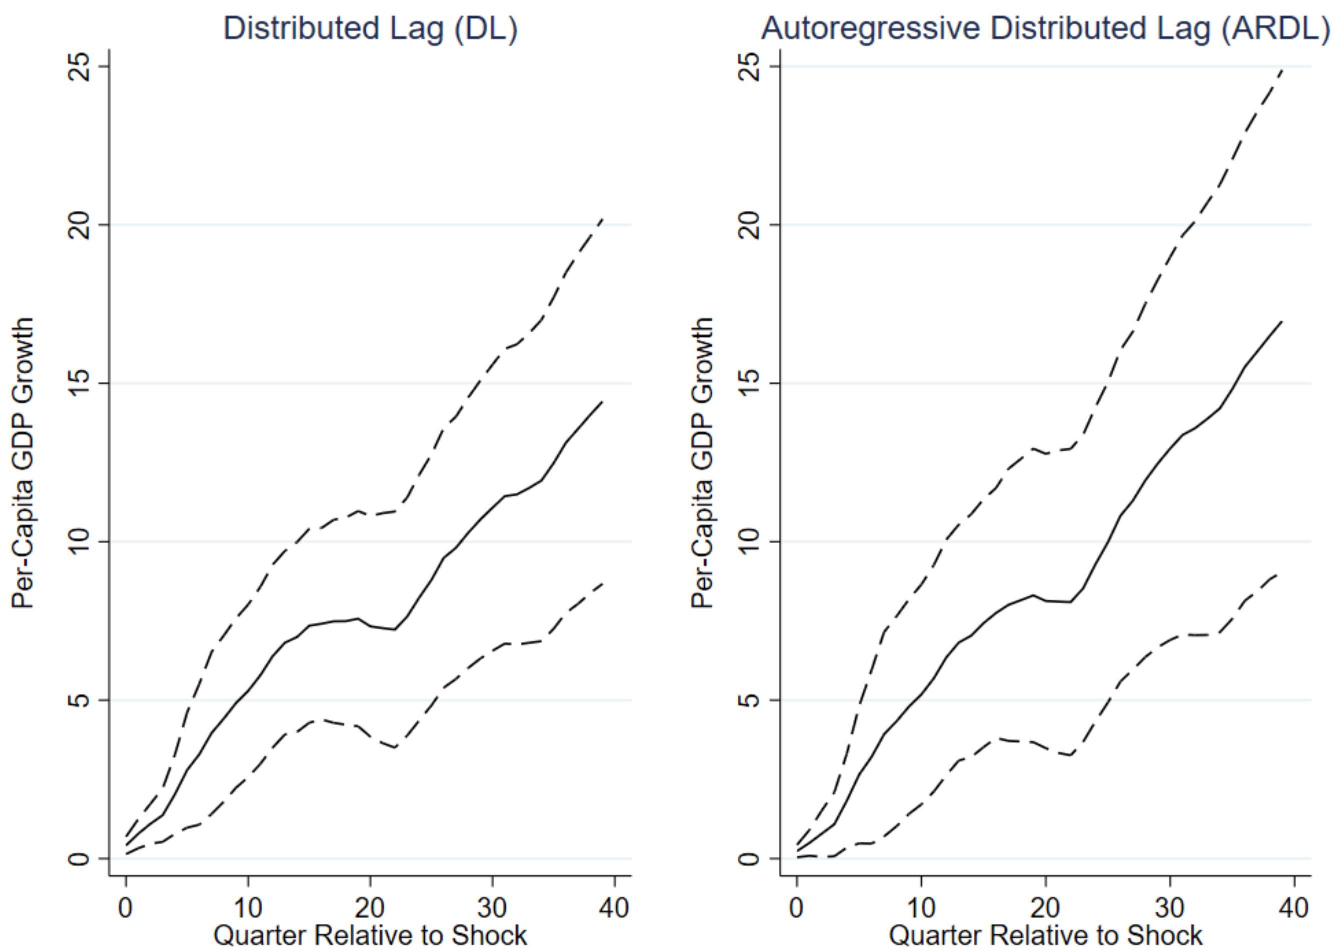

Note: The DL specification – as given by Equation (DL) – regresses quarter-over-quarter GDP growth on an indicator variable for the year-quarter in which a flat tax reform was adopted and 39 lags thereof, along with country and time fixed-effects. The indicator variable equals 1 in the year-quarter in which a flat tax reform was adopted, -1 in the year-quarter in which it was repealed, and 0 otherwise. The sample includes all 17 Eastern European and Central Asian flat-tax-adopting countries for which tax code and quarterly GDP data is available. The ARDL specification adds four autoregressive terms (lags of quarterly GDP growth) to the right-hand-side.

Figure 3: Effect of Flat Tax Reform on GDP Growth –Difference-in-Differences Specifications

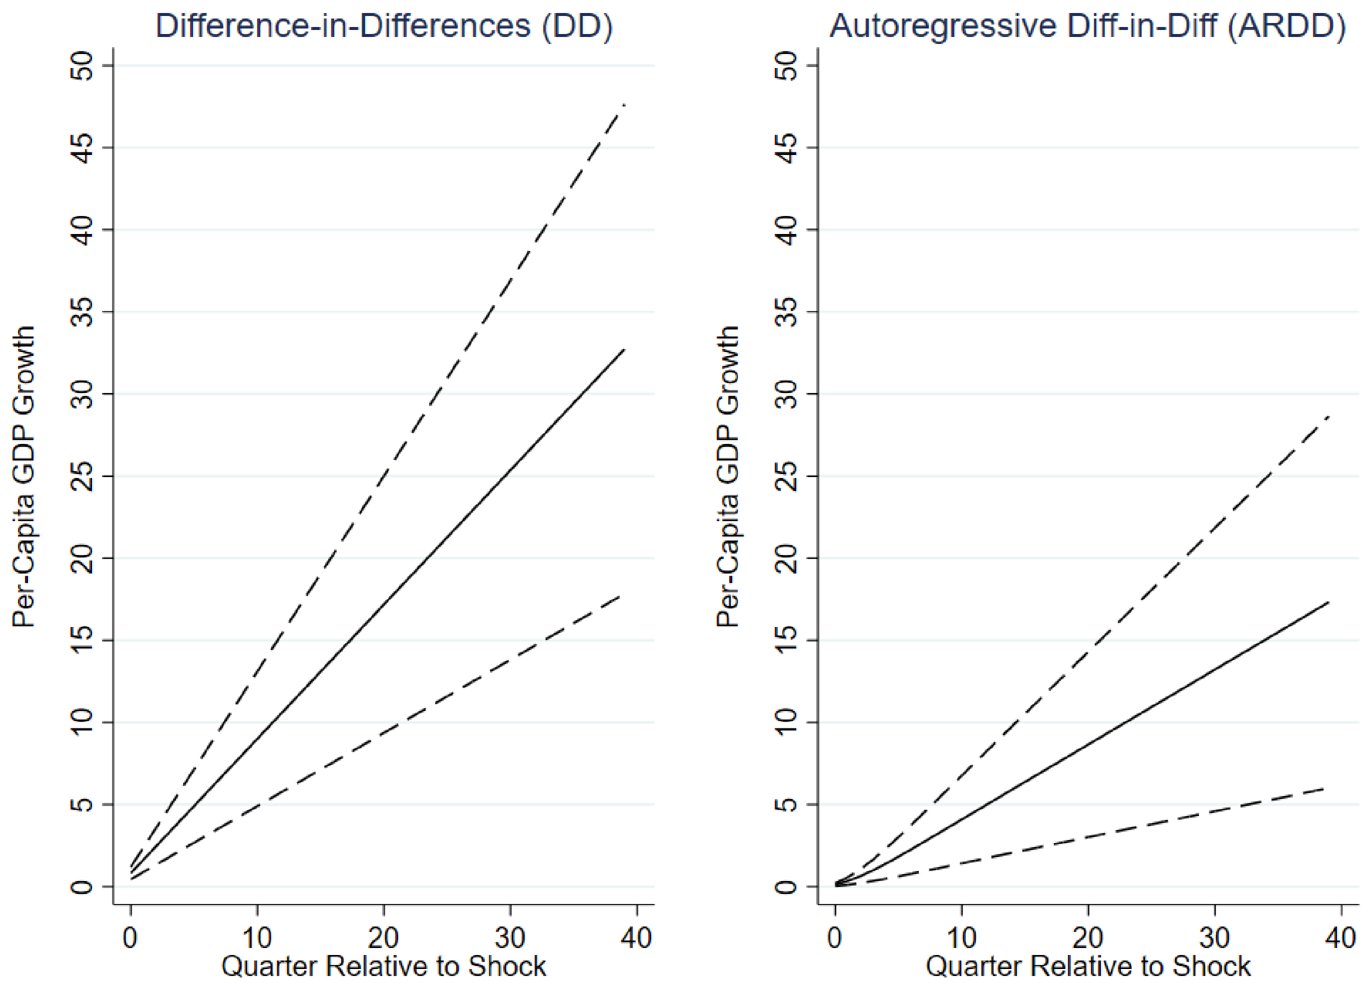

Note: The DD specification – as given by Equation (DD) – is a fairly standard difference-in-differences regression featuring a flat tax indicator variable and country and time fixed effects. The ARDD specification adds 4 lags of GDP growth to the DD specification.

Figure 4: Effect of Flat Tax Reform on GDP Growth – Event-Study Difference-in-Differences Specification

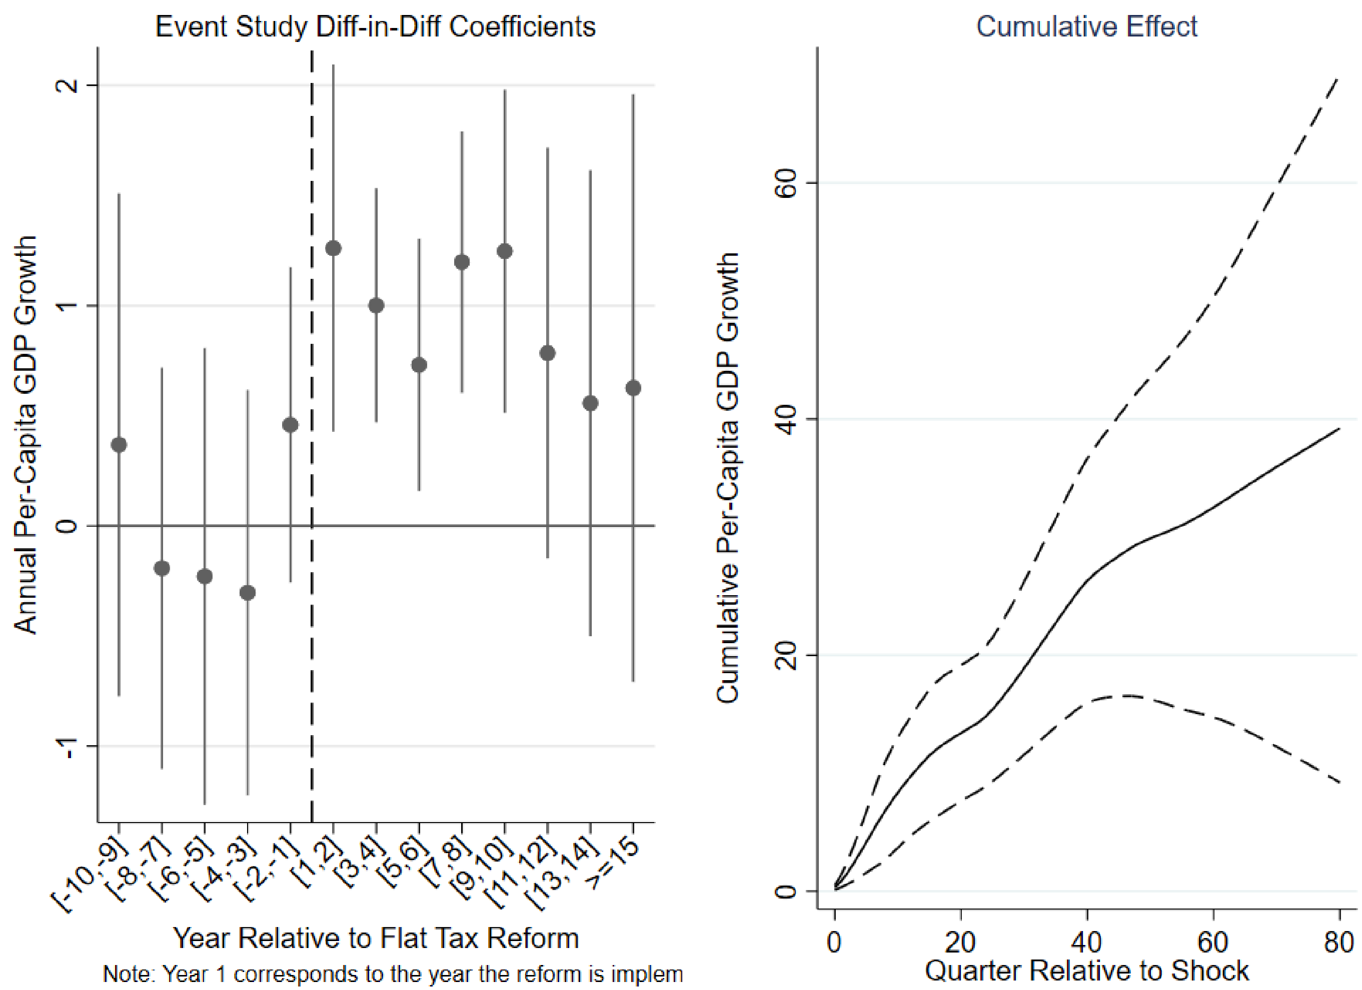

Note: This event study difference-in-differences specification is given by Equation (EventDD). The left panel shows the event-study coefficients and reveals the lack of any statistically significant pre-trends. It also shows that the flat tax reform only has statistically-significant effects on growth for the first decade, after which any additional effects on growth are not statistically distinguishable. This can also be seen in the right panel, which displays cumulative effects.

Figure 5: Goodman-Bacon Decomposition

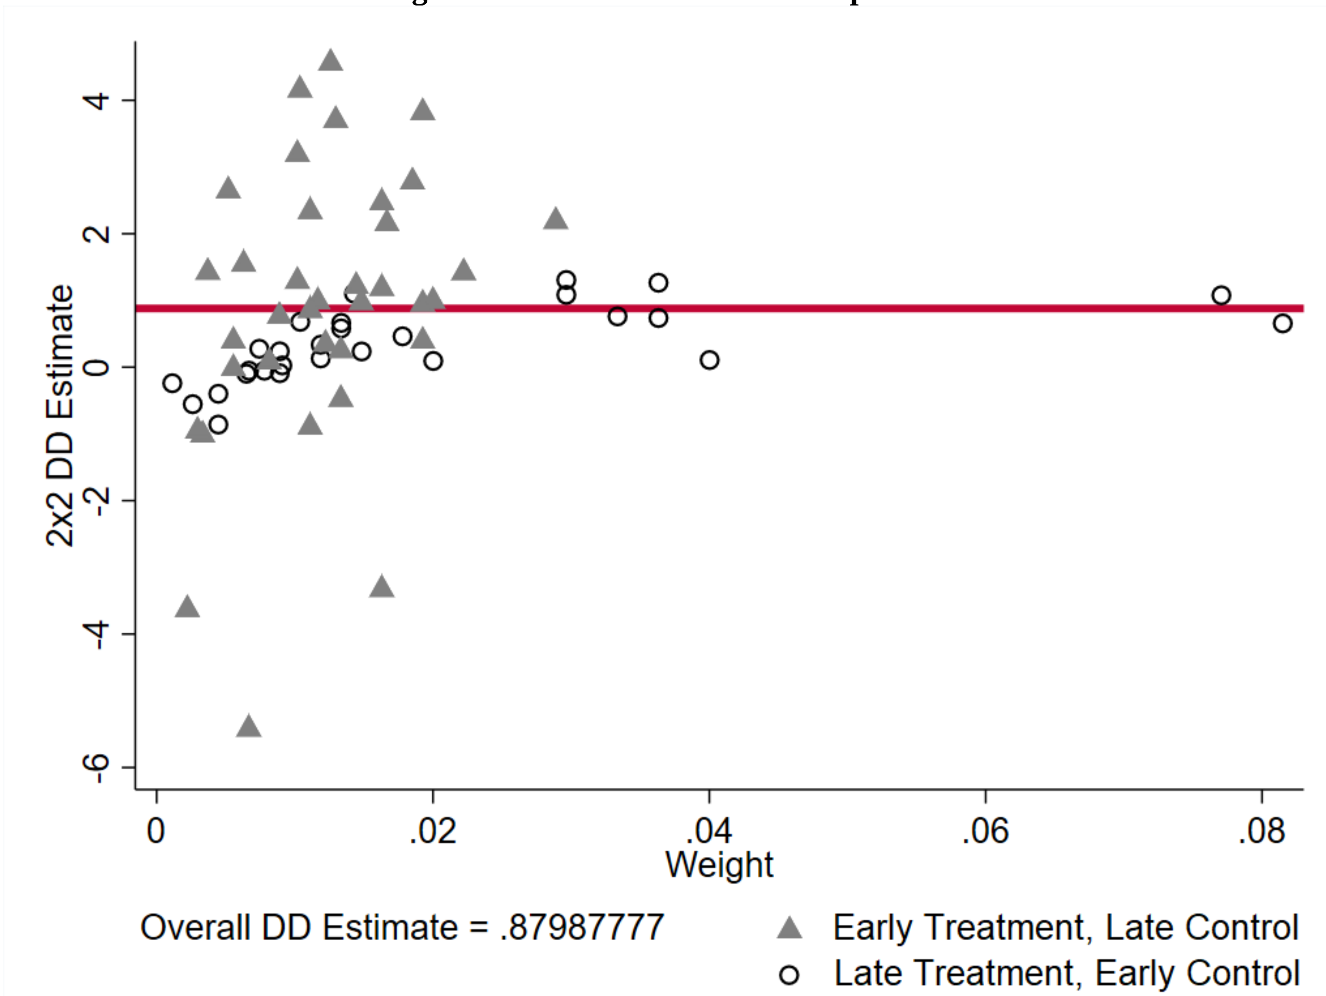

Note: This plot shows the results of a Goodman-Bacon (2021) decomposition intended to test the assumption of homogeneous treatment effects across timing of treatment. The treatment effect represented here is the quarterly effect on growth. The difference in average treatment effect across the two estimation types is not statistically significant.

Figure 6: Randomization Inference

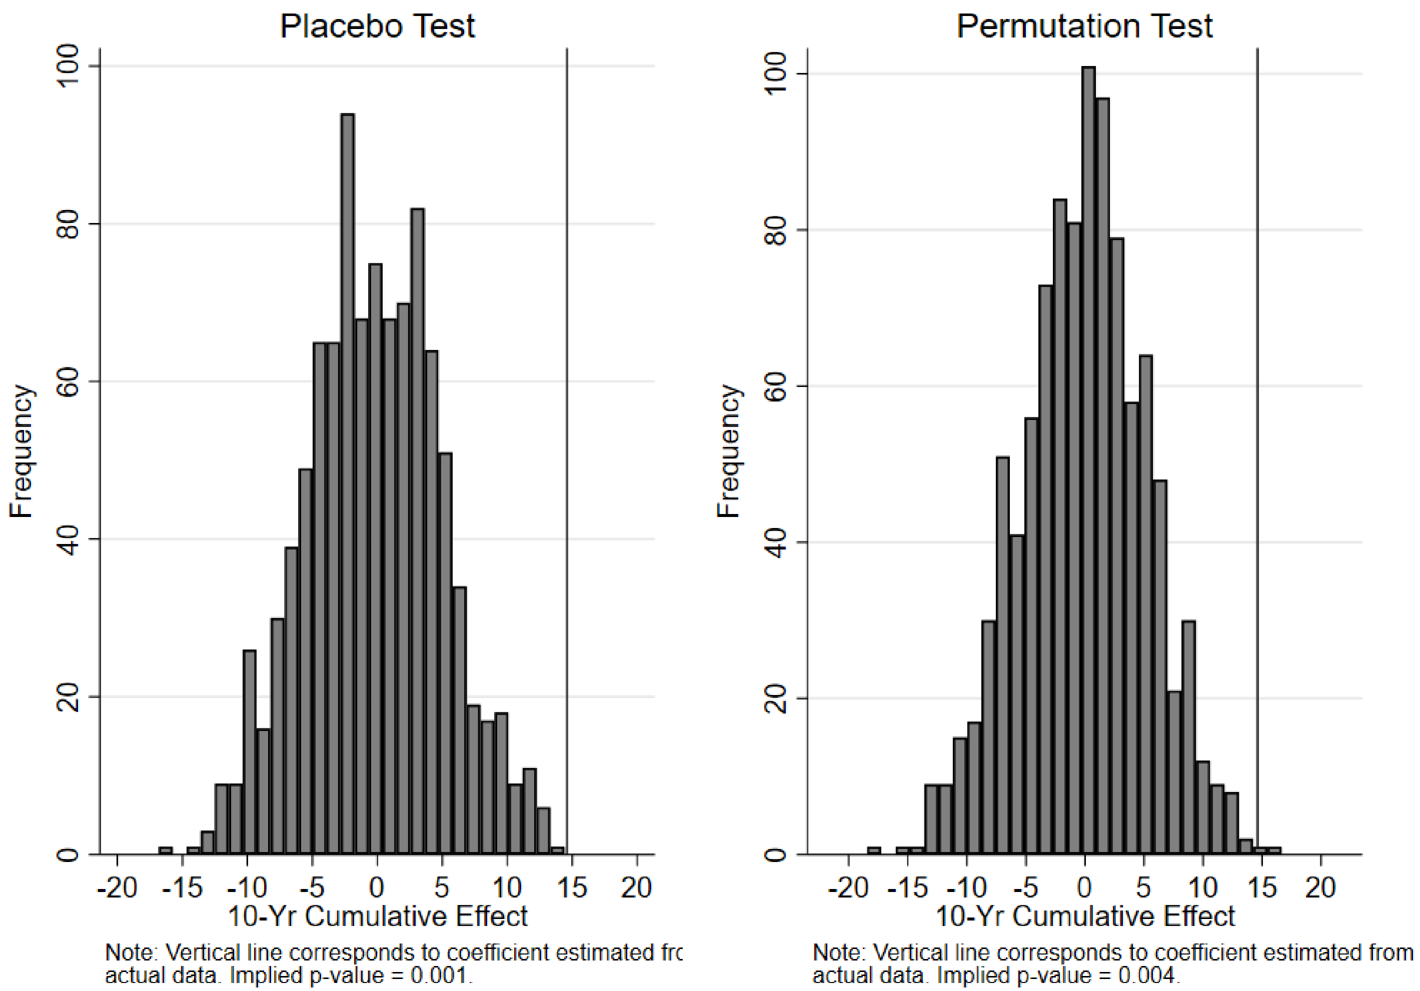

Note: The left panel displays the results of a 1000-repetition placebo test on the distributed-lag specification depicted in Figure 3, whereby each country in the panel is randomly assigned a placebo reform year and five are randomly assigned placebo repeal years. The right panel displays the results of a 1000-repetition permutation test on the same specification. The true 10-year effects are represented by the vertical lines.

Figure 7: Effects of AMTR and SDMTR Changes on GDP Growth – Distributed-Lag Specification

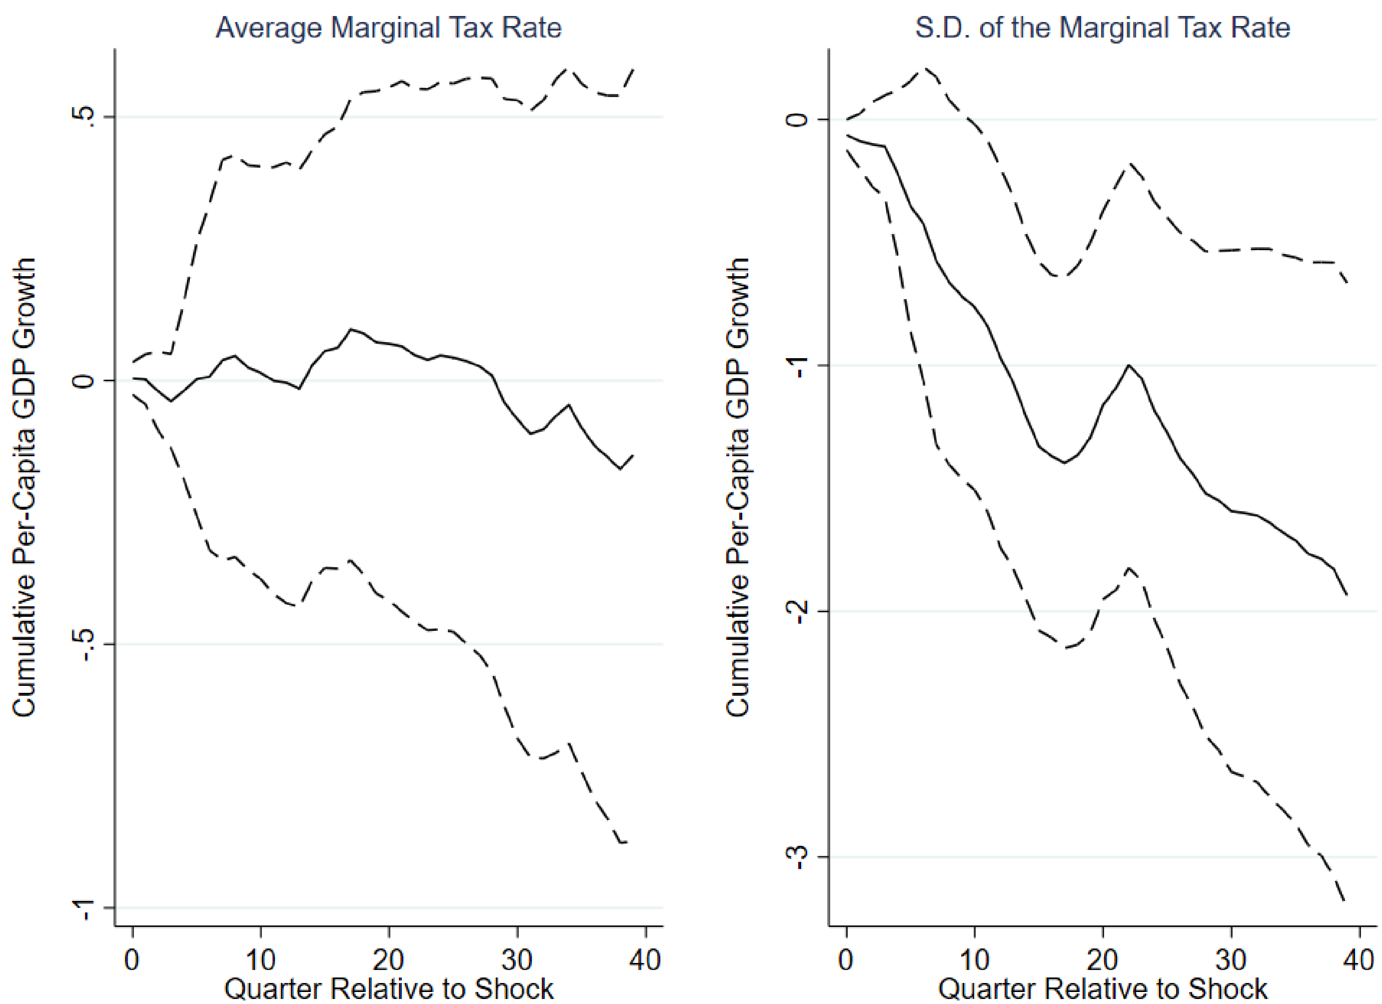

Note: This distributed-lag specification regresses quarter-over-quarter GDP growth regresses quarter-over-quarter GDP growth on a measure of the AMTR and SDMTR changes associated with the flat tax reforms and repeals along with 39 lags thereof, along with country and year-quarter fixed-effects.

Figure 8: Solow Model Example

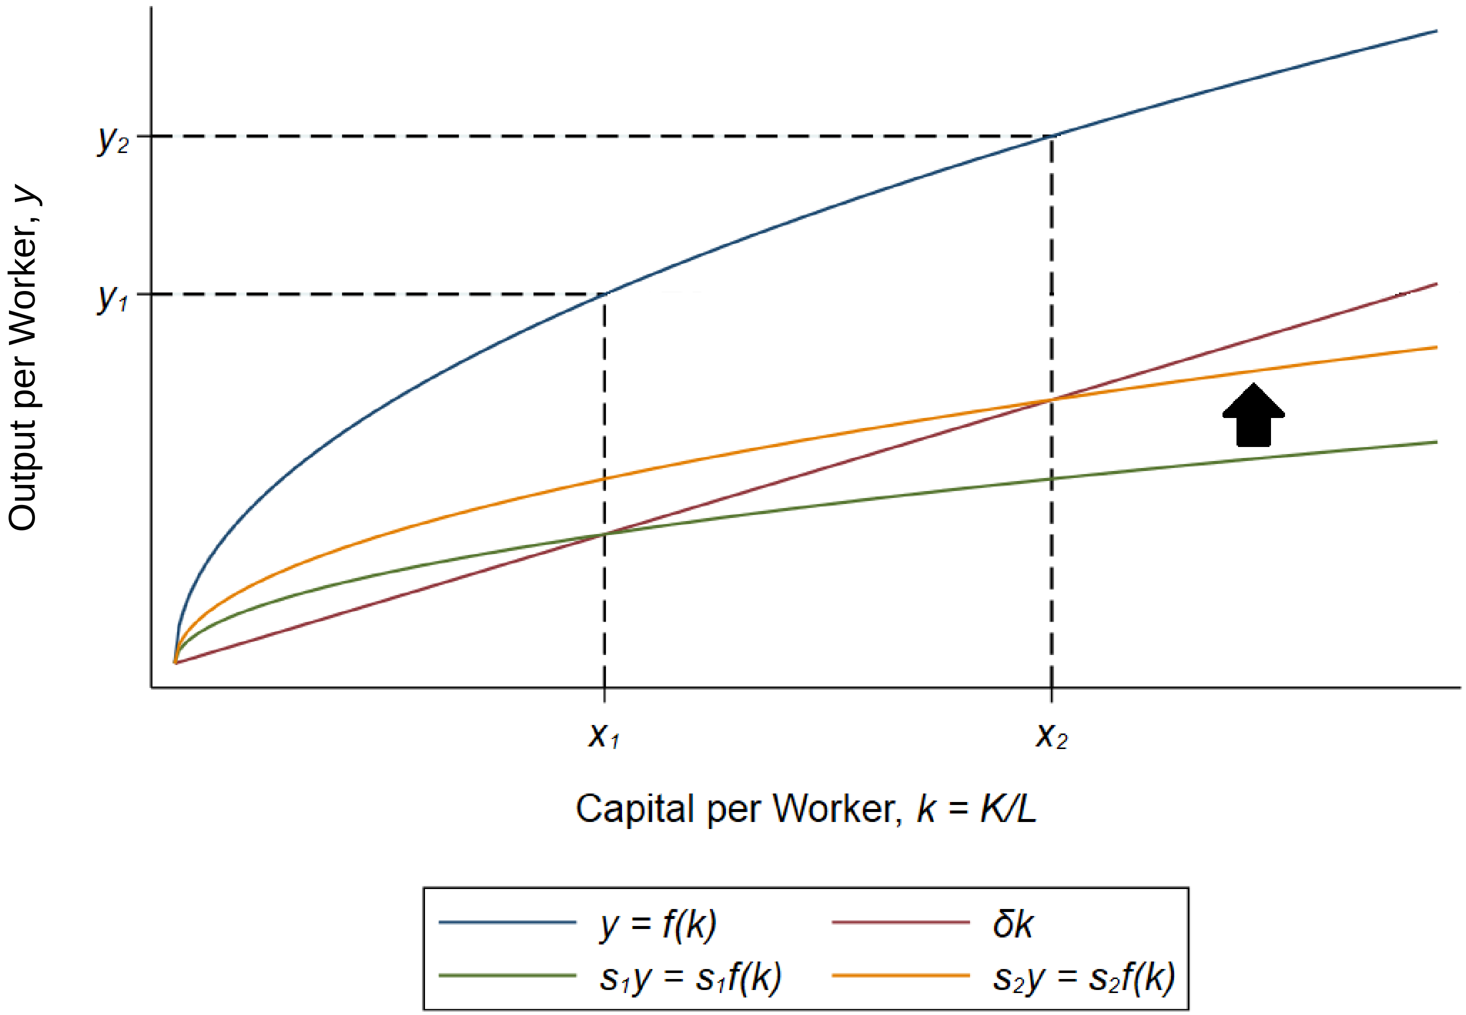

Note: An increase in saving/investment from s_1 to s_2 leads to an increase in the steady-state level of capital per-worker from x_1 to x_2. This is associated with an increase in steady-state output per-worker from y_1 to y_2. Since these are steady-state levels, there will be a transitionary period of adjustment from the old steady-state to the new steady-state, which will feature higher output growth.

Appendix A: Proof of Proposition

Observe that the individual can spend no more than his labor income wL – plus any investment income – over the course of his life. Thus, in a setting without taxation, his budget constraint would be given by $c_2 = (1 + r)S$, where $S \equiv wL - c_1$. However, in the present setting, the individual must pay tax on his income.

By definition, the individual saves and invests $S \equiv wL - c_1$ at the end of the first period. At the start of the second period, he receives investment income of rs and pays a total tax bill on his labor and investment income of $\tau(wL + rS)*(wL + rS) = \alpha(wL + rS) + \beta(wL + rS)^2$. Consequently, his budget constraint dictates period-2 consumption of

$$c_2 = (1 + r)S - \alpha(wL + rS) - \beta(wL + rS)^2.$$

Therefore, re-stating both c_1 and c_2 in terms of L and S yields a utility function in two variables:

$$u(c_1, L, c_2) = u\left(wL - S,\ \ L,\ \ (1+r)S - \alpha(wL + rS) - \beta(w^2L^2 + 2wLrS + r^2S^2)\right).$$

For simplicity, since utility is now a function of two endogenous variables, L and S, we can define

$$V(L,S) \equiv u\left(wL - S,\ \ L,\ \ (1+r)S - \alpha(wL + rS) - \beta(w^2L^2 + 2wLrS + r^2S^2)\right)$$

Differentiating the utility function with respect to L and S and setting these expressions equal to zero in order to obtain a maximum,

$$V_L = u_{c_1} \cdot w + u_L - u_{c_2} \cdot (\alpha w + 2\beta w^2 L + 2\beta wrS) = 0,$$

$$V_s = -u_{c_1} + u_{c_2} \cdot (1 + r - \alpha r - 2\beta r^2 S - 2\beta wLr) = 0.$$

Per the Implicit Function Theorem, we can learn how the optimal values of saving, S^*, and labor, L^*, vary as the tax parameters α and β vary. In particular,

$$\partial S^* / \partial \alpha = \frac{-V_{L\alpha}V_{LS} + V_{LL}V_{S\alpha}}{(V_{LS})^2 - V_{LL}V_{SS}},$$

$$\partial L^* / \partial \alpha = \frac{-V_{S\alpha}V_{LS} + V_{SS}V_{L\alpha}}{(V_{LS})^2 - V_{LL}V_{SS}}.$$

Note that the denominator of these expressions is unambiguously negative because the second derivative test dictates that $(V_{LS})^2 - V_{LL}V_{SS} < 0$ at the optimum values of L^* and S^*. With regard to the numerator,

$$
\begin{aligned}
&-V_{L\alpha}V_{LS}+V_{LL}V_{S\alpha} \\
&=\left(u_{c_1c_1}w+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)(1+r-\alpha r-2\beta r^2S-2\beta wrL)+u_{c_2}2\beta wr\right) \\
&\quad\left[u_{c_2c_2}\cdot(wL+rS)(\alpha w+2\beta w^2L+2\beta wrS)-u_{c_2}\cdot w\right] \\
&\qquad+\left(u_{c_1c_1}w^2+u_{LL}+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)^2-u_{c_2}\cdot 2\beta w^2\right) \\
&\qquad\left[-u_{c_2c_2}\cdot(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)-u_{c_2}\cdot r\right] \\
&=\left(u_{c_1c_1}w+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)(1+r-\alpha r-2\beta r^2S-2\beta wrL)+u_{c_2}2\beta wr\right) \\
&\quad\left[u_{c_2c_2}\cdot(wL+rS)(\alpha w+2\beta w^2L+2\beta wrS)-u_{c_2}\cdot w\right] \\
&\qquad+\left(u_{c_1c_1}w^2+u_{LL}+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)^2-u_{c_2}\cdot 2\beta w^2\right) \\
&\qquad\left[-u_{c_2c_2}\cdot(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)-u_{c_2}\cdot r\right] \\
&=\left(u_{c_1c_1}w+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)(1+r-\alpha r-2\beta r^2S-2\beta wrL)\right) \\
&\quad\left[u_{c_2c_2}\cdot(wL+rS)(\alpha w+2\beta w^2L+2\beta wrS)-u_{c_2}\cdot w\right] \\
&\qquad+\left(u_{c_1c_1}w^2+u_{LL}+u_{c_2c_2}\cdot(\alpha w+2\beta w^2L+2\beta wrS)^2\right) \\
&\qquad\left[-u_{c_2c_2}\cdot(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)-u_{c_2}\cdot r\right] \\
&\qquad\qquad+2\beta\begin{pmatrix} u_{c_2}u_{c_2c_2}\cdot wr(wL+rS)(\alpha w+2\beta w^2L+2\beta wrS) \\ +u_{c_2}u_{c_2c_2}\cdot w^2(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)\end{pmatrix}
\end{aligned}
$$

$(1 + r - \alpha r - 2\beta r(wL + rS)) > 0$ as a result of the assumption that the marginal tax rate never exceeds 1. Consequently, all terms in the above expression are unambiguously positive, except $-u_{c_2}\cdot r$. The expression is thus positive if

$$
\begin{aligned}
&-u_{c_2c_2}\cdot(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)-u_{c_2}\cdot r>0 \\
&\Leftrightarrow \frac{-u_{c_2c_2}c_2}{u_{c_2}}<\frac{rc_2}{(wL+rS)(1+r-\alpha r-2\beta r^2S-2\beta wLr)}
\end{aligned}
$$

This proves the first part of the proposition for saving, *s*. The proof for labor, *L*, is directly analogous.

Next, note that $\partial S^*/\partial\beta$ and $\partial L^*/\partial\beta$ are *not* the comparative statics of interest when it comes to determining how savings and labor change as tax progressivity changes while holding the average tax rate constant. Increasing β without modifying α will increase the average tax rate at any (positive) level of income.

So, consider an increase in β from β_1 to β_2. In order to keep an individual's overall tax rate constant, how must α change? For a given income level y_0,

$$\alpha_1 + \beta_1 y_0 = \alpha_2 + \beta_2 y_0 \Leftrightarrow \alpha_2 = \alpha_1 + (\beta_1 - \beta_2)y_0 .$$

Thus, if β increments by Δ, α must decrement by Δy_0. Consequently, in order to prove the proposition, it is necessary to show that the directional derivative

$$\partial s^* \big/ \partial\beta - (L + rS)\,\partial S^* \big/ \partial\alpha = \frac{-V_{LS}\left(V_{L\beta} - (wL + rS)V_{L\alpha}\right) + V_{LL}\left(V_{S\beta} - (wL + rS)V_{S\alpha}\right)}{(V_{LS})^2 - V_{LL}V_{SS}} < 0 .$$

As before, the denominator is unambiguously negative. Some simplification reveals that the numerator is unambiguously positive, yielding an expression that is indeed unambiguously negative overall. Beginning with the terms in the parentheses of the numerator,

$$\begin{aligned} V_{L\beta} - (wL + rS)V_{L\alpha} &= u_{c_2c_2} \cdot (wL + rS)^2(\alpha w + 2\beta w^2 L + 2\beta wrS) - u_{c_2} \cdot 2w(wL + rS) \\ &\quad - \left[u_{c_2c_2} \cdot (wL + rS)^2(\alpha w + 2\beta w^2 L + 2\beta wrS) - u_{c_2} \cdot w(wL + rS) \right] \\ &= -u_{c_2} \cdot w(wL + rS) \\ V_{S\beta} - (wL + rS)V_{S\alpha} &= -u_{c_2c_2} \cdot (wL + rS)^2(1 + r - \alpha r - 2\beta r^2 S - 2\beta wLr) - u_{c_2} \cdot 2r(wL + rS) \\ &\quad - \left[-u_{c_2c_2} \cdot (wL + rS)^2(1 + r - \alpha r - 2\beta r^2 S - 2\beta wLr) - u_{c_2} \cdot r(wL + rS) \right] \\ &= -u_{c_2} \cdot r(wL + rS) \end{aligned}$$

Next, analyzing the numerator as a whole,

$$\begin{aligned} & -V_{LS}\left(V_{L\beta} - (wL + rS)V_{L\alpha}\right) + V_{LL}\left(V_{S\beta} - (wL + rS)V_{S\alpha}\right) \\ &= \left(u_{c_1c_1} + u_{c_2c_2} \cdot (\alpha w + 2\beta w^2 L + 2\beta wrS)(1 + r - \alpha r - 2\beta r^2 S - 2\beta wLr) + u_{c_2} 2\beta wr\right)\left(-u_{c_2} \cdot w(wL + rS)\right) \\ &\quad + \left(u_{c_1c_1} \cdot w^2 + u_{LL} + u_{c_2c_2} \cdot (\alpha w + 2\beta w^2 L + 2\beta wrS)^2 - u_{c_2} \cdot 2\beta w^2\right)\left(-u_{c_2} \cdot r(wL + rS)\right) \\ &= \left(u_{c_1c_1} + u_{c_2c_2} \cdot (\alpha w + 2\beta w^2 L + 2\beta wrS)(1 + r - \alpha r - 2\beta r^2 S - 2\beta wLr)\right)\left(-u_{c_2} \cdot w(wL + rS)\right) \\ &\quad + \left(u_{c_1c_1} \cdot w^2 + u_{LL} + u_{c_2c_2} \cdot (\alpha w + 2\beta w^2 L + 2\beta wrS)^2\right)\left(-u_{c_2} \cdot r(wL + rS)\right) > 0 \end{aligned}$$

Therefore, as the progressivity of the tax schedule increases (holding constant the overall tax rate), saving/investment S declines unambiguously. The proof for labor L is directly analogous. This proves the proposition.

Appendix B: Measuring Progressivity

The Average Marginal Tax Rate (AMTR) is a measure that has been widely used in the macro-public finance literature. It dates back to Barro and Sahasakul (1983, 1986), providing a macro-level measure of the marginal tax rates faced by a typical unit of income in a given country. It is calculated as follows:

$$AMTR = \sum_b ShareIncome_b \cdot \tau_b$$

That is, the AMTR is a weighted average of the individual marginal tax rates, τ_b, in a country's tax schedule, where each τ_b is weighted by the share of total income accruing to individuals in the corresponding tax bracket, *b*. For instance, consider a country with two tax brackets, 20% below 1000 units of currency and 30% above 1000 units of currency. If half the population makes 800 units of currency in a year and the other half makes 1200 units, the average marginal tax rate is 25%. Even though the bulk of income in this hypothetical economy was taxed at the 20% rate, half of the populace faces the 30% on *any marginal income* that they earn. This is what the AMTR measures. As discussed by Barro and Sahasakul (and a multitude of more recent papers), because individuals respond to marginal rates rather than average rates in a whole range of economic decision-making, the AMTR is a more useful concept for macro-level examination of the response of investment, labor supply, etc. to various incentive changes.

The standard deviation of the marginal tax rate (SDMTR) is a useful extension of this concept that is amenable to measuring tax progressivity.

$$SDMTR = \sqrt{\sum_b ShareIncome_b \cdot (\tau_b - AMTR)^2}$$

Consider, for example, a pure flat tax system wherein every individual pays 20% on all income. In this case, because the MTR, τ_b, is equal to the AMTR throughout the tax schedule, the SDMTR will be precisely zero. The greater the commonality of deviations in the MTR from the AMTR (i.e., the higher the progressivity), the higher the value of the SDMTR.

In order to compute either the AMTR or the SDMTR, however, it is necessary to have a measure of $ShareIncome_b$ in each bracket *b*. This can be done using data on the tax codes of each of these countries, which I obtain from Ernst & Young and PricewaterhouseCoopers International Tax Guides, coupled with data from the World Income Inequality Database (WIID) on share of total income by decile of population. Evidence suggests that 99% of the income distribution is well-approximated by a lognormal distribution (Clementi and Gallegati 2005).

The WIID data can thus be fit to a lognormal distribution, with mean

$$\mu = \frac{\sum_{j=1}^{10} \ln y_j}{10}, \text{ where } y_j = \frac{s_j \cdot GNI}{POP/10}$$

and variance

$$\sigma^2 = \frac{\sum_{j=1}^{10} (y_j - \mu)^2}{10-1},$$

where s_j is the share of income accruing to population decile *j* (ordered by income), *GNI* is gross national income, and *POP* is national population. y_j, then, is average gross income in population decile *j*. With this information, one can compute the proportion *P*(*y*) of the population with incomes below *y*:

$$P(y) = \Phi\left(\frac{\ln y - \mu}{\sigma}\right),$$

where Φ denotes the CDF of the normal distribution. Furthermore, the share *S*(*y*) of income accruing to people with incomes below *y* is given by the Lorenz Curve of the lognormal distribution:

$$S(y) = \Phi\left(\frac{\ln y - \mu - \sigma^2}{\sigma}\right)$$

Next, equipped with the tax code data, one can define

$$y^b \equiv \frac{D + max\_inc^{b-1}}{1-\tau_{ep}},$$

where *D* denotes the standard deduction, max_inc^b denotes the upper bound of tax bracket *b*, and τ_{ep} denotes the payroll tax paid by employees. Note that the denominator of y^b is instead simply equal to 1 if the employee's contribution to payroll tax is not deductible from income tax in the country in question. max_inc^0 is defined to be 0. Thus y^b is a measure of the actual maximum gross income that will be exposed to bracket *b*, after accounting for deductions[13]. In practice, virtually all "flat" taxes effectively have two brackets: one at 0% for income below the amount of the personal exemption and one at the official flat rate for income above the amount of the

[13] In general, the tax code of these countries is substantially less complex than that of the United States. While some of the countries do grant additional deductions, e.g., for children or for being married, for the purposes of this macro-level exercise, I abstract from these considerations. The final result is unlikely to be sensitive to these factors. If, for example, I disregard the deductibility of payroll taxes, the relevant regression specifications (those in Table 6 which make use of the AMTR/SDMTR variables) are virtually unchanged.

personal exemption. Consequently, AMTR and SDMTR can be computed as follows:

$$AMTR = \sum_b (S(y^b) - S(y^{b-1})) \cdot \tau_b$$

$$SDMTR = \sqrt{\sum_b (S(y^b) - S(y^{b-1})) \cdot (\tau_b - AMTR)^2} \, .$$