

Putting the ‘Finance’ into ‘Public Finance’: A Theory of Capital Gains Taxation*

Mark Aguiar [†]

Benjamin Moll [‡]

Florian Scheuer [§]

First version: May 2024

This version: June 2024

[\[latest version\]](#)

Abstract

Standard optimal capital tax theory abstracts from modeling asset prices, making it unsuitable for thinking about capital gains and wealth taxation. We study optimal redistributive taxation in an environment with asset price changes, adopting the modern finance view that asset prices fluctuate not only because of changing cash flows, but also due to other factors (“discount rates”). We show that the optimal tax base (i) generally differs from the case with constant asset prices, and (ii) depends on the sources of asset-price changes. Whenever asset prices fluctuate, and are not exclusively driven by cash flow changes, taxes must target realized trades and generally involve a combination of realization-based capital gains and dividend taxes. This result stands in contrast to the classic Haig-Simons comprehensive income tax concept, as well as recent proposals for wealth or accrual-based capital gains taxes.

*We thank Mike Golosov, Camille Landais, Dmitry Mukhin, Daniel Reck, Dejanir Silva, Joel Slemrod, Iván Werning and many seminar participants for helpful discussions. Erik Beyer and Seyed Hosseini Maasoum provided excellent research assistance. Benjamin Moll acknowledges support from the Leverhulme Trust and the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 865227). Florian Scheuer acknowledges support from ERC Starting Grant No. 757721 and SNSF Consolidator Grant No. 213673.

[†]Princeton University

[‡]London School of Economics

[§]University of Zurich

“Many of the distortions associated with the present system of capital gains taxation result from its deviation from the Haig-Simons approach. These deviations may have historical explanations but their persistence is hard to rationalize from an economic perspective.” (Auerbach, 1989)

The treatment of capital gains due to changing asset prices lies at the heart of many debates regarding the taxation of capital income and wealth. While capital gains are typically taxed on realization (i.e. asset sale) in practice, a long tradition in public finance going back to [von Schanz \(1896\)](#), [Haig \(1921\)](#) and [Simons \(1938\)](#) advocates for taxing capital gains on accrual. This idea has recently made its way into policy proposals, including by the Biden administration.¹ In the United States, such tax policies would invariably end up in the Supreme Court which has, for more than a hundred years, heard cases about, but never conclusively ruled on, whether unrealized gains constitute income.² Finally, debates about wealth taxation often end up being about the desirability (and practicality) of taxing wealth changes due to asset-price movements, simply because such unrealized capital gains typically dwarf ordinary saving and income flows for top wealth holders.

The existing public finance literature on optimal capital taxation abstracts from explicitly modelling changing asset prices, and therefore provides no guidance in these debates.³ Our paper aims to fill this gap in the literature by “putting the ‘finance’ into ‘public finance’.” That is, we study optimal redistributive taxation in the presence of asset price fluctuations. Importantly, we do so adopting the view of the modern finance literature that asset prices change not only in response to changing cash flows but also due to changes in discount rates ([Campbell and Shiller, 1988](#)). In this dichotomy, “discount rates” simply means any sources of asset price changes other than current and expected future cash flows. Empirically, asset prices move too much to be accounted for by changing cash flows alone, both at high frequencies and over longer time horizons.⁴

Our main contribution is to show that optimal redistributive taxes (i) generally differ from the case with constant asset prices, and (ii) depend on the sources of asset-price changes. Whenever asset prices fluctuate beyond movements due to cash flow changes, taxes must target realized trades, and generally involve a combination of realization-based capital gains

¹U.S. Office of Management and Budget (2022), U.S. Department of the Treasury (2022), [Saez et al. \(2021\)](#), [Zucman \(2024\)](#), and [The Economist \(2024\)](#). [Leiserson and Yagan \(2021\)](#) calculate that the 400 wealthiest U.S. families paid an average tax rate of only 8.2% in the years 2010 to 2018 by including unrealized capital gains in the tax base.

²Most recently *Moore v. United States* in 2024 but going back to *Eisner v. Macomber* in 1920. The key open question is whether unrealized gains constitute income under the 16th Amendment of the U.S. constitution. See [Fox and Liscow \(2024\)](#) for a useful summary of the legal arguments and the U.S. Supreme Court’s position.

³See, for example, the classic contributions of [Atkinson and Stiglitz \(1976\)](#), [Chamley \(1986\)](#), and [Judd \(1985\)](#). Like us, [Piketty et al. \(2023\)](#) lament the absence of asset-price effects from this literature but nevertheless proceed without explicitly modeling them. Interestingly, [Lucas \(1990\)](#) begins his review of the literature with a discussion of capital gains taxation: “When I left graduate school, in 1963, I believed that the single most desirable change in the U.S. tax structure would be the taxation of capital gains as ordinary income. I now believe that neither capital gains nor any of the income from capital should be taxed at all. My earlier view was based on what I viewed as the best available economic analysis, but of course I think my current view is based on better analysis.” However, also Lucas does not explicitly consider changing asset prices (though see the discussion of Lucas’ quote in Section 1).

⁴See for example [Shiller \(1981\)](#), [Campbell and Shiller \(1988\)](#), [Cochrane \(2011\)](#), [Greenwald et al. \(2019\)](#) and [van Binsbergen \(2020\)](#) as well as the secular increase in many measures of price-dividend ratios and the decline in real interest rates. While this is the conventional view, others have argued that fluctuations in cash flows are first order. Our reading of this debate is that it is imperative to understand the tax implications of both sources.

and dividend taxes. The reason is that, holding constant cash flows, asset-price increases redistribute toward asset sellers who realize capital gains, away from asset purchasers who pay a higher price for a given dividend stream while not directly affecting those who do not trade. Optimal redistributive taxation must take this dynamic into account, as well as accounting for any changes in relative income due to cash flows changes.

Taxes that are optimal in environments with constant asset prices may cease to be optimal, or even change in counterintuitive ways, when asset prices fluctuate. While a wealth tax may be optimal with constant asset prices, its progressivity needs to change whenever asset prices move and optimal taxation may even prescribe tax cuts for the wealthiest when asset prices rise. Taxing unrealized capital gains is optimal only in restrictive knife-edge cases, so that our results also stand in contrast to the classic Haig-Simons comprehensive income tax concept.⁵

We begin our study of redistributive taxation with changing asset prices in an environment that is purposely simple: a deterministic small open economy in which a large number of investors trade one asset with an exogenously given asset price and homogeneous asset returns across investors. Investors begin with heterogeneous endowments of the asset and face different income profiles. The small open economy assumption allows us to study the implications of fluctuations in asset prices and cash flows on the income distribution in a transparent manner. Later in the paper, we study various extensions, including leveraged borrowing and general equilibrium, and show that the results from this simple setting generalize.

We are interested in how the optimal tax system responds to changes in asset prices. As a first step, we assume the government has access to type-specific lump-sum taxes and characterize the set of first-best tax schedules that trace out the Pareto frontier. This benchmark is useful as it generates a clear distinction in how taxes respond to changes in discount rates versus cash flows. We then show that the principles observed in the first-best problem are present in a second-best allocation in which the government is restricted to distortive taxes à la [Mirrlees \(1971\)](#), so that the classic tradeoff between redistribution and efficiency arises. While the first-best is clearly not realistic and implies extreme predictions about optimal tax *rates*, it turns out to be instructive about the optimal tax *base*, i.e., what taxes should condition on depending on the sources of asset price changes. These results then generalize in a natural way to more interesting second-best tax systems.

To explain our findings, it is useful to consider the standard definition for an asset's return

$$R_{t+1} = \frac{D_{t+1} + p_{t+1}}{p_t}, \quad (1)$$

where p_t denotes the asset's price and D_t its cash flow, i.e. the return equals dividend yield plus capital gain. While the existing literature features either a constant asset price or small movements of asset returns (see the literature discussion below), we instead allow both cash flows and asset prices, and hence returns, to change in flexible ways. In particular, we allow for changes in asset returns R_t that are independent of changes in cash flows D_t , i.e. discount

⁵Immediately after the quote at the beginning of this introduction, [Auerbach \(1989\)](#) adds: "It is therefore disappointing and puzzling that the debate about capital gains taxes continues to focus almost exclusively on tax rates rather than on tax structure." We wholeheartedly agree.

rate changes.

Consider the following thought experiment. The economy is initially in a steady state with asset price \bar{p} , dividend \bar{D} , and associated asset return \bar{R} . In this initial steady state, there is a tax system in place that optimally redistributes across investors according to some Pareto weights. Now suppose that, at time $t = 0$, the time paths of asset prices, cash flows, and hence returns change to some alternative time paths $\{p_t, D_t, R_t\}$. In particular, asset prices p_t , may increase because expected future cash flows increase or for other reasons, i.e. discount rate changes. The question we are after is: how should the tax system respond to these changes?

A useful stepping stone for answering this question is the idea of “Slutsky compensation,” defined as the change in the investor’s budget that keeps the initial consumption bundle affordable at the new prices and dividends. We show that this compensation generally requires conditioning on realized trades: when asset prices rise, sellers benefit and hence need to be taxed whereas buyers lose and hence need to be compensated. Of course, dividend income changes are also compensated or taxed. Building on the Slutsky-compensation logic, optimal taxation in the first-best problem is straightforward: just like Slutsky compensation, it taxes sellers, compensates buyers, and taxes dividend income changes. Importantly, optimal lump-sum taxes generally target realized trades rather than asset holdings.

There are two useful polar special cases. In the first special case, the time path of asset prices $\{p_t\}$ changes while cash flows remain at the initial steady state \bar{D} . Because asset returns going forward $\{R_t\}$ change, this case corresponds to asset price changes driven entirely by discount rates. In the second special case, asset prices and cash flows $\{p_t, D_t\}$ instead change proportionately and in such a way that the asset return remains at the initial steady state \bar{R} . This corresponds to asset price changes driven entirely by cash flows.

We show that, in the first special case with changing discount rates, optimal lump-sum taxes depend *only* on investors’ realized trades (purchases and sales) as well as the price changes relative to steady state, and are independent of investors’ asset holdings. Intuitively, rising asset prices benefit sellers, who are therefore taxed, and hurt buyers, who are therefore subsidized. In contrast, in the second special case with changing cash flows, optimal lump-sum taxes target the investor’s individual wealth gain due to the asset price change, so that it is asset holdings rather than transactions that matter. However, this is a knife-edge result: whenever asset-price changes are not exclusively driven by cash flow changes, optimal lump-sum taxes target realized trades as well.

While our formula for optimal redistributive taxes is reminiscent of realization-based capital gains taxation systems observed in practice, it also differ in important ways. For example, optimal taxes (i) not only tax sellers but also compensate buyers who experience “purchasing losses” when prices rise; (ii) they compensate realized capital *losses* and tax “purchasing gains” when prices fall; (iii) they tax *net* rather than *gross* transactions (selling and re-investing at the same price should not incur a tax liability) and (iv) the capital gain or loss is typically calculated relative to a basis that differs from the historical basis at which the investor purchased the asset. Finally, in this first-best case with non-distortive taxation, our formula prescribes an extreme tax rate of 100%, i.e. the government taxes away all realized capital gains in their

entirety and uses the proceeds to compensate the losers from rising asset prices.

Turning to second-best tax systems à la [Mirrlees \(1971\)](#), our results regarding the optimal tax *base* carry over from the first-best analysis in a natural way. In this environment, only a subset of investor choices can be taxed, for example asset sales, consumption, or savings. We show that the incentive compatibility constraints share some key similarities across the various settings. Our main interest remains how the second-best tax schedule changes when asset prices rise. We show that the second-best tax schedule monotonically increases as a function of trading gains, albeit with a slope less than in the first-best. This is intuitive – taxing asset sellers in response to a price increase achieves a better distribution of income, like in the first-best, but also distorts saving behavior. If end-of-period wealth is taxed rather than sales, optimal taxes may become *less* progressive when asset prices rise. Intuitively, if those holding the asset at the end of the period are net *purchasers*, they should be subsidized rather than taxed (relative to the baseline tax schedule). While this example is extreme, it illustrates why the fluctuating market value of investors' asset holdings is a problematic target for redistributive taxes.

We also show how these results depend on the household's inter-temporal elasticity of substitution (IES). Given that the distortion generated by second-best taxation is on saving behavior, a lower IES brings us closer to the first-best tax schedule. In particular, we prove that, in the limit as the IES goes to zero, the second-best tax schedule achieves the first-best allocation. This establishes that the insights from the first-best tax system are not based on a knife-edge scenario, but extend qualitatively to more realistic environments.

Finally, we consider various extensions, including risk and borrowing, bequests, return heterogeneity, and general equilibrium considerations. While these extensions change our optimal tax formula in natural and instructive ways, the key findings emphasized in this introduction remain unchanged. Specifically, in contrast to the classic Haig-Simons comprehensive income tax concept as well as proposals for wealth or accrual-based capital gains taxes, optimal redistributive taxes generally target realized trades. An extension in which investors trade the same asset but receive heterogeneous cash flows makes clear just how knife-edge the case is in which wealth or accrual-based capital gains taxes are optimal. Specifically, suppose that cash flows increase but only for a subset of the investors, thereby increasing the equilibrium asset price. Then the remaining investors experience an increase in the asset price but without a corresponding increase in cash flows, exactly as if there had been a lower discount rate, with sellers benefiting and buyers losing. Therefore, taxes on unrealized capital gains or wealth are no longer optimal even when asset prices are driven entirely by changing cash flows.

Our extensions with borrowing and heterogeneous returns also allow us to speak to an issue that has received attention in the popular debate: wealthy individuals borrowing against appreciating assets rather than selling them, often aiming to take advantage of the “stepped-up basis” for bequeathed assets as part of a “buy, borrow, die” tax avoidance strategy. Our results suggest that basis step-up should be eliminated, thereby also eliminating the viability of this strategy. Absent the stepped-up basis, the wealthy would still benefit from borrowing against high-return assets with lower-interest loans. But exploiting such return differences is a feature of *any* levered investment strategy and should not be considered tax avoidance.

Literature. Our paper contributes to the theoretical public finance literature studying the optimal taxation of capital income and wealth. Apart from the classic contributions we already mentioned, see the references in Section 1.4 as well as the surveys by [Golosov et al. \(2007\)](#), [Banks and Diamond \(2010\)](#), [Bastani and Waldenstrom \(2020\)](#), [Stantcheva \(2020\)](#) and [Scheuer and Slemrod \(2021\)](#). To differentiate our paper from this literature, it is again useful to consider the expression for an asset’s return (1).

The existing literature features either a constant asset price (and hence no capital gains or losses whatsoever) or works with variants of the neoclassical growth model. In the growth model, asset-return movements are typically small, reflecting the disappointing asset-pricing properties of the standard real business cycle model. Our analysis instead allows for flexible changes in asset returns R_t that are independent of changes in cash flows D_t , i.e. discount rate changes. Within the environments it has considered, the literature has shown that taxing asset *holdings* may be optimal, for example by means of a wealth tax. Our paper instead shows that such taxes are problematic whenever asset prices fluctuate and are not exclusively driven by cash flow changes. In all such cases, taxes must target realized trades, and generally involve a combination of realization-based capital gains and dividend taxes.

In line with our argument that it is essential to “put the ‘finance’ into ‘public finance’”, a growing positive literature has documented an important role for asset-price and interest-rate changes in driving wealth inequality and has studied the macroeconomic and distributional implications of these trends (e.g. [Bonnet et al., 2014](#); [Rognlie, 2015](#); [Kuhn et al., 2020](#); [Gomez, 2016](#); [Wolff, 2022](#); [Gomez and Gouin-Bonenfant, 2020](#); [Cioffi, 2021](#); [Catherine et al., 2020, 2024](#); [Greenwald et al., 2021](#); [Moll, 2020](#); [Martínez-Toledano, 2022](#); [Fagereng et al., 2023](#)). The logic of our results is closely related to [Moll \(2020\)](#) and [Fagereng et al. \(2023\)](#) who study the welfare-relevant redistributive effects of changing asset prices. Our paper contributes to this literature by instead studying the normative implications of changing asset prices, specifically their implications for optimal capital taxation.

There is also an empirical literature studying behavioral responses, specifically of asset sales, to capital gains taxation aiming to estimate the relevant elasticities.⁶ While our paper tackles optimal distortive taxation à la [Mirrlees \(1971\)](#) only in a stylized two-period model which is not suitable for making quantitative predictions about optimal tax rates, such elasticities will of course be key inputs in more serious quantitative work.

While the modern theoretical public finance literature provides no guidance on the optimal tax treatment of capital gains, there is an older, mostly verbal or graphical literature that anticipates some of the key ideas in our paper, in particular [Paish \(1940\)](#), [Kaldor \(1955\)](#) and [Whalley \(1979\)](#), which were partly reactions to the work of [Haig \(1921\)](#) and [Simons \(1938\)](#) that developed the eponymous income concept.

⁶See for example the handbook chapter by [Poterba \(2002\)](#), [Feldstein et al. \(1980\)](#) and [Agersnap and Zidar \(2021\)](#). The literature has paid particular attention to the “lock-in effect” of realization-based capital gains taxation and has proposed various solutions to eliminate this incentive to postpone selling appreciated assets, see e.g. [Vickrey \(1939\)](#), [Auerbach \(1989\)](#), and [Auerbach \(1991\)](#). There are also theoretical and quantitative studies of behavioral responses including the lock-in effect (e.g. [Constantinides, 1983](#); [Chari et al., 2005](#); [Smith and Miller, 2023](#)).

Roadmap. Section 1 spells out our baseline environment. Section 2 focuses on a special case with two time periods that we use in much of our subsequent analysis. Section 3 studies the first-best allocation assuming that the government has access to type-specific lump-sum taxes. In contrast, Section 4 incorporates the classic tradeoff between redistribution and efficiency in a second-best world with distortive taxation. Section 5 shows how our findings regarding optimal taxes carry over to the multi-period model of Section 1. Section 6 considers various extensions, including risk and borrowing, bequests, return heterogeneity, and general equilibrium considerations. Section 7 concludes.

1 Baseline model

We begin by spelling out our baseline environment, which we then use to study optimal redistributive taxation in the following. This baseline model is kept purposely simple; in particular, we study a deterministic small open economy in which a large number of investors trade one asset with an exogenously given asset price and homogeneous asset returns across investors. Section 6 considers various extensions, including general equilibrium, return heterogeneity, risk and borrowing, and intergenerational considerations. For now, we omit taxes from the analysis, which we will introduce in Section 3.

1.1 Investors

Time is discrete with time periods $t = 0, 1, \dots, T$, where T may be finite or infinite. There is a continuum of heterogeneous investors indexed by their type $\theta \in [\underline{\theta}, \bar{\theta}]$, which is distributed in the population according to the cumulative distribution function $F(\theta)$. Investors have preferences over consumption in the T time periods, (c_0, \dots, c_T) , captured by the utility function $U(c_0, \dots, c_T)$, which is strictly increasing and strictly concave in all its arguments. Investors receive type-specific exogenous income flows $\{y_t(\theta)\}_{t=0}^T$. They can transfer income across periods by saving in an asset that pays a deterministic dividend stream $\{D_t\}_{t=0}^T$ and trades at an exogenously given price $\{p_t\}_{t=0}^T$. Investors have an initial asset endowment $k_0(\theta)$ and decide how many assets $\{k_{t+1}(\theta)\}_{t=0}^T$ to carry into the next period. There is no short-selling constraint, and hence $k_{t+1}(\theta)$ may be negative or positive. We could also allow for additional assets, for example a risk-free bond, but these would be perfect substitutes in our deterministic environment and would therefore not affect our analysis. We consider the more interesting case with multiple assets that are imperfect substitutes in Section 6.

The problem of an investor of type θ is to maximize her welfare

$$U(\theta) = \max_{\{c_t(\theta), k_{t+1}(\theta)\}_{t=0}^T} U(c_0(\theta), \dots, c_T(\theta)),$$

subject to the sequence of budget constraints:

$$c_t(\theta) + p_t(k_{t+1}(\theta) - k_t(\theta)) = y_t(\theta) + D_t k_t(\theta) \quad \forall t \geq 0, \quad (2)$$

with $k_0(\theta)$ given and where we impose $p_T = 0$ in the final period.

The only fundamental underlying heterogeneity across investors is in initial asset holdings $k_0(\theta)$ as well as income over time $\{y_t(\theta)\}_{t=0}^T$. This heterogeneity generates gains from trade, with natural buyers and sellers of the asset. Figure 1 depicts an example with two time periods in which high θ types have lower initial income $y_0(\theta)$ but higher future income $y_1(\theta)$ and relatively similar initial asset holdings $k_0(\theta)$. In this example, low θ types will be buyers of the asset (savers with $c_0(\theta) < y_0(\theta)$) whereas high θ types will be sellers (effective borrowers with $c_0(\theta) > y_0(\theta)$).




Figure 1: Two-period example of heterogeneity in initial assets and incomes over time

Notes: In this two-period example, low θ types who have high initial income $y_0(\theta)$ and low future income $y_1(\theta)$ will be buyers of the asset (savers) whereas high θ types for whom the opposite is true will be sellers (borrowers).

The asset return is given by equation (1) in the introduction: $R_{t+1} = (D_{t+1} + p_{t+1})/p_t$, i.e. the return equals dividend yield plus capital gain. The investor's intertemporal optimality condition (Euler equation) is

$$U_{c_t}(c_0(\theta), \dots, c_T(\theta)) = R_{t+1} U_{c_{t+1}}(c_0(\theta), \dots, c_T(\theta)), \quad (3)$$

where $U_{c_t} \equiv \partial U / \partial c_t$ for $t = 0, \dots, T$. For future reference, it is also useful to define the cumulative return between time t and $t + s$ for $s \geq 1$:⁷

$$R_{t \rightarrow t+s} \equiv R_{t+1} \cdot R_{t+2} \cdots R_{t+s}, \quad (4)$$

Haig-Simons income. The budget constraint (2) states that “consumption plus saving equals income.” An equivalent way of writing this accounting identity adds unrealized capital gains $(p_t - p_{t-1})k_t(\theta)$ on both sides, thus changing the definitions of saving and income (consumption is unchanged):

$$c_t(\theta) + \underbrace{p_t k_{t+1}(\theta) - p_{t-1} k_t(\theta)}_{\text{change in wealth}} = \underbrace{y_t(\theta) + D_t k_t(\theta) + (p_t - p_{t-1}) k_t(\theta)}_{\text{Haig-Simons income}} \quad \forall t \geq 0.$$

⁷To obtain compact present-value summations, we will denote no discounting by $R_{t \rightarrow t}$, i.e. $R_{t \rightarrow t} \equiv 1$.

Formulation (2) features disposable income, whereas this formulation features “Haig-Simons income” which includes unrealized capital gains (Haig, 1921; Simons, 1938). Defining the market value of wealth⁸ $a_t(\theta) \equiv p_{t-1}k_t(\theta)$ and the net return including capital gains $r_t = (D_t + p_t - p_{t-1})/p_{t-1} = R_t - 1$, Haig-Simons income also equals $y_t(\theta) + r_t a_t(\theta)$, i.e. income including total capital income. Similarly, adding $a_t(\theta)$ on both sides of the budget constraint, yields the standard

$$c_t(\theta) + a_{t+1}(\theta) = y_t(\theta) + R_t a_t(\theta) \quad \forall t \geq 0, \quad (5)$$

with $a_0(\theta) = p_{-1}k_0(\theta)$ given.

A useful utility function. We have not imposed any assumptions on the utility function U besides it being strictly increasing and strictly concave. But for many of our results it will be useful to specialize this utility function to:

$$U(c_0, \dots, c_T) = G(C(c_0, \dots, c_T)) \quad \text{where} \\ C(c_0, \dots, c_T) = \left(\sum_{t=0}^T \beta^t c_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{and} \quad G(C) = \frac{C^{1-\gamma}}{1-\gamma} \quad \text{with} \quad \sigma, \gamma > 0. \quad (6)$$

Here, C is a composite commodity in which β is the discount factor used to discount consumption in the second time period and σ is the intertemporal elasticity of substitution. The parameter γ governs curvature over this composite commodity. Note that this utility is a monotone transformation of the more standard intertemporally separable utility function $\sum_{t=0}^{\infty} \beta^t u(c_t)$, with $u(c) = \frac{c^{1-1/\sigma}}{1-1/\sigma}$. As a result, the investor’s Euler equation (3) specializes to the standard Euler equation

$$c_t(\theta)^{-1/\sigma} = \beta R_{t+1} c_{t+1}(\theta)^{-1/\sigma}.$$

The reason for working with this monotone transformation is that some of our results involve taking the limit of the utility function as the intertemporal elasticity of substitution σ goes to zero, which is ill-defined for the standard specification.⁹ In contrast, the utility function (6) satisfies $U(c_0, \dots, c_T) \rightarrow G(\min\{c_0, \dots, c_T\})$ as $\sigma \rightarrow 0$, i.e., it converges to a (monotone transformation of a) Leontief utility function, as expected.

1.2 Aggregate economy

The economy’s aggregate resource constraint is found by simply aggregating investors’ budget constraints (2) across individuals. To this end, we use the convention to denote aggregate variables by capital letters, for example

$$C_t = \int c_t(\theta) dF(\theta), \quad K_t = \int k_t(\theta) dF(\theta),$$

⁸The mismatching time subscripts in $a_t(\theta) \equiv p_{t-1}k_t(\theta)$ are solely due to our notational convention which uses $k_t(\theta)$ to denote asset holdings *at the beginning* of period t . Alternatively, using $k_t(\theta)$ to denote asset holdings *at the end* of period t , (2) becomes $c_t(\theta) + p_t(k_t(\theta) - k_{t-1}(\theta)) = y_t(\theta) + D_t k_{t-1}(\theta)$ so that wealth is $a_t(\theta) \equiv p_t k_t(\theta)$.

⁹For example $\frac{c^{1-1/\sigma}}{1-1/\sigma} \rightarrow 0$ as $\sigma \rightarrow 0$ for all $c > 1$ (the numerator converges to zero and the denominator to $-\infty$).

and so on. With this notation, the aggregate resource constraint is

$$p_t K_{t+1} + C_t = (p_t + D_t)K_t + Y_t \quad (7)$$

for all t . As already noted, our benchmark analysis focuses on a small open economy with an exogenously given time path for asset prices and dividends $\{p_t, D_t\}_{t=0}^T$. Hence, the economy's aggregate asset holdings at time $t + 1$, K_{t+1} , may differ from those at t , K_t , as the economy as a whole may be a net buyer or net seller of K . Let $X_t \equiv K_t - K_{t+1}$ denote aggregate net sales. In Section 6.1, we alternatively consider a closed economy general equilibrium version of the model in which the asset that investors save in is in fixed supply, so that asset sales or purchases are zero in the aggregate: for every seller, there is a buyer.

1.3 Sources of asset-price changes

Our interest is in the taxation of gains or losses due to changes in asset prices. The asset-pricing literature draws an important distinction between different sources of asset price changes, in particular distinguishing between asset discount rates and cash flows. In this dichotomy, "discount rates" simply means any sources of asset price changes other than current and expected cash flows. Using a log-linear accounting decomposition of observed asset price changes due to [Campbell and Shiller \(1988\)](#), much of this literature has found that discount rate shocks account for most of asset price fluctuations.¹⁰ Other studies have argued that fluctuations in cash flows are first order.¹¹ Our reading of this debate is that it is imperative to understand the tax implications of both sources.

To this end, consider the equation for the asset return (1) but adopt the perspective of the asset-pricing literature to treat required asset returns or discount rates $\{R_t\}$ as a primitive and prices $\{p_t\}$ as an outcome. One way of thinking about this is that, in equilibrium models, it is typically the asset returns $\{R_t\}$ that are pinned down which, in turn, determine asset prices.¹² Rearranging (1) as $p_t = R_{t+1}^{-1}(D_{t+1} + p_{t+1})$ and iterating forward using $R_{t \rightarrow t+s}$ defined in (4) yields

$$p_t = \sum_{s=1}^{T-t} \frac{D_{t+s}}{R_{t \rightarrow t+s}}, \quad (8)$$

i.e. the asset price equals the present-discounted value of dividends. There are thus two sources of rising asset prices: rising dividends $\{D_{t+s}\}$ and declining discount rates $\{R_{t+s}\}$. [Campbell and Shiller \(1988\)](#) provide a log-linear accounting decomposition of (8) that makes this precise.

Our analysis below considers optimal taxation when asset prices change relative to some baseline. One useful example of such a baseline is the Gordon growth model ([Gordon and](#)

¹⁰See [Campbell \(2018, Section 5.3.1\)](#) for an expository treatment and derivation of the Campbell-Shiller decomposition.

¹¹See for example, [Larrain and Yogo \(2008\)](#) and [Atkeson et al. \(2024\)](#).

¹²For example, in the two-period general equilibrium version of our model in section 6.1, the investors' Euler equation together with market clearing pins down the equilibrium return as $R^* = \frac{1}{\beta} \left(\frac{Y_0}{Y_1 + DK_0} \right)^{-1/\sigma}$. This then pins down the equilibrium asset price as $p^* = D/R^* = \beta \left(\frac{Y_0}{Y_1 + DK_0} \right)^{1/\sigma}$ which is the expression in equation (34).

Shapiro, 1956) in which $T = \infty$, discount rates are constant $R_{t+1} = \bar{R}$ and cash flows grow at a constant rate $\bar{D}_{t+1}/\bar{D}_t = G > 1$ so that

$$\bar{p}_t = \frac{\bar{D}_{t+1}}{\bar{R} - G}. \quad (9)$$

In the Gordon growth model, asset price growth is entirely due to dividend growth and the price-dividend ratio is constant. Another example of a baseline is simply a steady state with constant dividend \bar{D} , return \bar{R} and asset price \bar{p} (the Gordon growth model with $G = 1$).

In our simple partial equilibrium model with exogenous prices and dividends $\{p_t, D_t\}$, we can then consider the following two extremal experiments to isolate discount rates and cash flows as the two drivers of asset price changes.

1. Changes in the path of the asset price $\{p_t\}$ holding dividends $\{D_t\}$ fixed. This corresponds to asset price changes driven entirely by discount rate changes (the asset return $\{R_{t+1}\}$ changes).
2. Changes in the path of the asset price $\{p_t\}$ that are accompanied by proportional changes in dividends $\{D_t\}$ such that the asset return $\{R_{t+1}\}$ stays constant.¹³

Both of these cases are the opposite extremes of the general case, with arbitrary changes in $\{p_t\}$ and $\{D_t\}$, which corresponds to asset price changes driven by a mixture of dividend and discount rate changes. One of our main arguments in Section 3 is that special case 2, in which the asset return stays constant because asset price increases are accompanied by proportional rising dividends or, equivalently, the asset price change is driven entirely by cash flows, is an extreme knife-edge case with some special results that do not carry over to any deviation from this configuration. Furthermore, as we show in Section 6, in a richer model with return heterogeneity, these results no longer apply even when asset prices are driven exclusively by dividend changes.

1.4 Comparison to setups studied in the capital taxation literature

Before proceeding, we briefly connect our setup to other models in the existing literature on optimal capital taxation. These make different assumptions on the determination of asset prices and dividends $\{p_t, D_t\}_{t=0}^T$, and hence returns $\{R_t\}_{t=0}^T$.

Partial equilibrium models with constant return. This is the special case in which $R_t = \bar{R}$ for all t . The most obvious way of generating this result is to assume that $p_t = \bar{p}$ and $D_t = \bar{D}$ for all t . Alternatively, prices and dividends could grow at the same constant rate. This captures models of capital taxation with a linear savings technology, such as the finite-horizon models based on Atkinson and Stiglitz (1976) (e.g. Saez, 2002; Scheuer and Wolitzky, 2016; Hellwig and Werquin, 2024; Ferey et al., 2024), some of the new dynamic public finance literature (surveyed in Golosov et al. (2007)), or infinite-horizon partial equilibrium models such as Saez and Stantcheva (2018).

¹³Section 5 spells out the conditions for asset-price and dividend changes to keep returns constant.

Neoclassical growth model. Starting with [Chamley \(1986\)](#), many papers have studied optimal capital taxation in variants of the growth model. Denote by $\sum_{t=0}^{\infty} \beta^t U(C_t)$ the preferences of the representative consumer and by $f(K_t, A_t L_t)$ the constant-returns technology for producing output, where C_t is consumption, K_t is capital, A_t is productivity, and L_t is labor with inelastic supply $L_t = 1$. How to map the growth model into our setup depends on the particular decentralization. In all decentralizations, the asset return is $R_{t+1} = f_K(K_{t+1}, A_{t+1}) + 1 - \delta$ and this asset return equals the relevant discount rate (this is the standard Euler equation):

$$R_{t+1} = \frac{1}{\beta} \frac{U'(C_t)}{U'(C_{t+1})}. \quad (10)$$

Furthermore, the unit price of capital (relative to consumption) equals one because the consumption good can be converted into investment one-for-one.¹⁴

In contrast, dividends and asset prices $\{D_t, p_t\}_{t=0}^{\infty}$ differ according to the particular decentralization and asset-market structure. Our asset could be physical units of capital, in which case the asset price is fixed at $p_t = 1$ and there are *no* capital gains or losses. Alternatively, it could be shares in the representative firm which are in unit fixed supply, the typical assumption in the literature studying asset pricing in production economies (e.g. [Jermann, 1998](#)). The cash flows D_t are then firm profits net of investment and the asset price equals the firm's capital stock $p_t = K_{t+1}$ (see [Appendix A.1](#)) so that there *are* capital gains and losses. An interesting case is that of a balanced growth path (BGP) with productivity growth $A_{t+1}/A_t = G > 1$ and isoelastic preferences $U'(C) = C^{-1/\sigma}$. On this BGP, the asset return is constant and pinned down from $\bar{R} = (1/\beta)G^{1/\sigma}$ but it consists of both a dividend yield and a capital gains component:

$$\frac{D_{t+1}}{p_t} = \bar{R} - G, \quad \frac{p_{t+1}}{p_t} = G.$$

Capital gains are driven entirely by growing cash flows and the price-dividend ratio is constant as in the Gordon growth model. The correct notion of capital income is the sum of dividend income plus (unrealized) capital gains. Chamley's result is that the long-run tax rate on this combined capital income should be zero (which makes sense of the Lucas quote in footnote 3).

However, even in this decentralization, it is still true that the price per unit of capital equals one. Furthermore, movements in the asset return R_{t+1} – e.g. along the transition or in response to movements in productivity – are quantitatively small, analogous to the disappointing asset-pricing properties of the real business cycle model.¹⁵

Growth models with heterogeneous households. Going back to [Judd \(1985\)](#), many contributions have studied capital taxation in growth models with heterogeneous households or entrepreneurs (see [Werning \(2007\)](#), [Shourideh \(2012\)](#), [Farhi et al. \(2012\)](#), [Straub and Wern-](#)

¹⁴We discuss this property in more detail in [Appendix A.1](#) where we also discuss how to break it.

¹⁵The difficulty with explaining asset returns in real-business cycle models is closely connected to the assumption that the consumption good can be converted into investment one-for-one. For example, [Jermann \(1998\)](#) writes that “in the standard one-sector model agents can easily alter their production plans to reduce fluctuations in consumption. This suggests that the frictionless and instantaneous adjustment of the capital stock is a major weakness in this framework.”

ing (2020), Benhabib and Szöke (2021) and Guvenen et al. (2023, 2024) for recent examples). Despite the (often) rich heterogeneity, the backbone of all of these papers is the neoclassical growth model so the discussion in the preceding paragraph still applies.

Our setup. The setups studied in the existing literature feature either constant asset prices (and hence no capital gains whatsoever) or small movements of asset returns (discount rates) and a constant unit price of capital. We instead study optimal taxation with exogenous time paths for asset prices and dividends $\{p_t, D_t\}_{t=0}^T$ and hence returns $\{R_t\}_{t=0}^T$. This simplification allows us to take on board the modern finance view that asset prices change not only because of changing cash flows but also due discount rates. While our baseline analysis is therefore silent on the ultimate fundamental drivers of asset prices (preferences and technology), Section 6 considers extensions featuring general equilibrium and risk premia and shows that our findings remain valid in these richer environments.

2 Two time periods

We now use the special case with two time periods, $t = 0, 1$, to illustrate the redistributive effects of asset-price changes. Much of our analysis of optimal taxation in subsequent sections similarly uses this two-period case before we show later that the findings carry over to the multi-period case in a natural fashion.

2.1 The investor's problem

With two time periods $t = 0, 1$, the investor's problem is to maximize utility $U(c_0(\theta), c_1(\theta))$ subject to the following budget constraints:

$$c_0(\theta) + p(k_1(\theta) - k_0(\theta)) = y_0(\theta), \quad (11)$$

$$c_1(\theta) = y_1(\theta) + Dk_1(\theta). \quad (12)$$

Given that the world ends after time period 1, the asset price $p_1 = 0$. For simplicity, we also assume that the asset pays no dividend in the first period $D_0 = 0$. Given this, we then drop the time-subscripts on p_0 and D_1 to ease notation. The asset's returns in the two time periods are given by

$$R_0 \equiv \frac{p}{p_{-1}} \quad \text{and} \quad R_1 \equiv \frac{D}{p}, \quad (13)$$

which is the standard expression (1) with $D_0 = p_1 = 0$ and where p_{-1} denotes a constant baseline price.

As will become clear, a useful reformulation of the investors' problem is in terms of asset sales

$$x(\theta) \equiv k_0(\theta) - k_1(\theta), \quad (14)$$

where $x > 0$ represents sales and $x < 0$ purchases of k . Using this definition, the investors'

problem becomes

$$\begin{aligned}
U(\theta) &= \max_{\{c_0(\theta), c_1(\theta), x(\theta)\}} U(c_0(\theta), c_1(\theta)) \quad \text{s.t.} \\
c_0(\theta) &= y_0(\theta) + px(\theta) \\
c_1(\theta) &= y_1(\theta) + D(k_0(\theta) - x(\theta))
\end{aligned} \tag{15}$$

The $t = 0$ budget constraint states that consumption $c_0(\theta)$ equals exogenous income $y_0(\theta)$ plus revenue from asset sales $px(\theta)$. The $t = 1$ budget constraint states that consumption $c_1(\theta)$ equals exogenous income $y_1(\theta)$ plus capital income $D(k_0(\theta) - x(\theta))$ consisting of the dividend payments D on the assets brought forward to period 1, $k_0(\theta) - x(\theta)$.

It is sometimes useful to combine the two period budget constraints in (15) into a present-value budget constraint:

$$c_0(\theta) + \frac{p}{D}c_1(\theta) = y_0(\theta) + \frac{p}{D}y_1(\theta) + pk_0(\theta), \tag{16}$$

This constraint states that the present-discounted value of consumption (discounted at the asset return $R = D/p$ defined in (13)) equals the present-discounted value of income plus initial wealth. This constraint can also be aggregated across all investors to yield

$$C_0 + \frac{p}{D}C_1 \leq Y_0 + \frac{p}{D}Y_1 + pK_0, \tag{17}$$

which is the economy's aggregate resource constraint in this partial equilibrium model.

Sources of asset-price changes and our two special cases. In the two-period model, the discussion of cash flows and discount rates as drivers of asset-price changes in Section 1.3 is particularly simple. Treating the required return R_1 and dividends D as the primitives in (13), the two-period version of (8) is simply $p = D/R_1$. Special Case 1 (only discount rate changes) is thus the case in which the asset price p changes holding dividends D fixed. On the opposite extreme, Special Case 2 (only cash flow changes) is the case in which both p and D change proportionately such that the asset return $R_1 = D/p$ stays constant.

2.2 An Envelope Condition

The goal of our paper is to study optimal redistributive taxation with changing asset prices; specifically, to study the optimal response of the tax system to rising asset prices. As a warm up, it is useful to first consider a simpler question: what are the redistributive effects of rising asset prices and cash flows, i.e. who wins and who loses as a result of these changes? Consider small deviations of the asset price dp and dividends dD . Following Moll (2020) and Fagereng et al. (2023), we use the envelope theorem to differentiate the value function $V(\theta)$ of investors defined in (15) to obtain

$$dU(\theta) = U_{c_0}(\theta) \left(x(\theta) dp + \frac{p}{D}k_1(\theta) dD \right). \tag{18}$$

The most interesting case is our Special Case 1: asset price changes but holding cashflows fixed, $dp > 0$ but $dD = 0$. In this case, the effect of a rise in p is given by marginal utility times the extent to which this rise relaxes the budget constraint at $t = 0$, namely asset sales $x(\theta)$ times the price change dp . Intuitively, a rising asset price benefits individuals who plan to sell the asset (i.e., $x(\theta) > 0$) and hurts individuals who plan to buy the asset (i.e., $x(\theta) < 0$). In particular, to first order, it does not affect individuals who do not plan to trade (i.e., $x(\theta) = 0$): for those individuals, the increasing asset price is merely a “paper gain” with no corresponding effect on consumption and thus welfare. To first order, only asset *transactions* matter whereas asset *holdings* do not.

When dividends rise as well $dD > 0$ —as in Special Case 2 and intermediate cases—the term involving $k_1(\theta) dD$ becomes non-zero. The intuition is that rising dividends directly benefit asset holders ($k_1(\theta) > 0$). However, it remains true that the direct welfare effect of the asset-price change dp itself depends only on asset transactions $x(\theta)$.

3 First best

We are interested in how the optimal tax system responds to changes in asset prices. As a first step, we will answer this question assuming that the government has access to type-specific lump-sum taxes (first best). While this is not realistic and implies extreme predictions about optimal tax *rates* (indeed, we will consider more realistic, second-best tax systems in Section 4), it turns out to be instructive about the optimal tax *base*, i.e., what taxes should condition on depending on the sources of asset price changes.

3.1 First-best consumption allocation

For a given asset price p and dividend D , any Pareto efficient allocation $\{c_0^*(\theta), c_1^*(\theta)\}$ solves

$$\max_{\{c_0(\theta), c_1(\theta)\}} \int \omega(\theta) U(c_0(\theta), c_1(\theta)) dF(\theta) \quad (19)$$

subject to the aggregate resource constraint (17), where $\omega(\theta)$ is the Pareto weight on an investor of type θ . While this planning problem can be solved for general preferences, the solution is particularly simple under the assumption that preferences are given by (6). In this case, the planner simply assigns to each investor θ a time-invariant share of (optimally chosen) aggregate consumption C_t^* :

$$c_t^*(\theta) = \Omega(\theta) C_t^*, \quad t = 0, 1, \quad (20)$$

where the share $\Omega(\theta)$ satisfies

$$\Omega(\theta) \equiv \frac{\omega(\theta)^{1/\gamma}}{\int \omega(\theta')^{1/\gamma} dF(\theta')}. \quad (21)$$

The optimal allocation can be implemented in a decentralized equilibrium when the government is able to redistribute with type-specific lump-sum taxes $T_0(\theta)$ in period 0 and $T_1(\theta)$

in period 1. The investors' budget constraints (11) and (12) become

$$c_0(\theta) = y_0(\theta) + px(\theta) - T_0(\theta) \quad (22)$$

$$c_1(\theta) = y_1(\theta) + D(k_0(\theta) - x(\theta)) - T_1(\theta) \quad (23)$$

We impose the government budget constraints

$$\int T_0(\theta)dF(\theta) = \int T_1(\theta)dF(\theta) = 0,$$

which implies, without loss, that the government does not own assets itself.

To back out the optimal taxes from the optimal consumption allocation $\{c_0^*(\theta), c_1^*(\theta)\}$, we can use the budget constraints (22) and (23). Given the ability of an investor to move resources inter-temporally, $T_0(\theta)$ and $T_1(\theta)$ are not separately pinned down and we require a normalization. One example, which we pick in the following, is to set $T_1(\theta) = 0$. Then the second-period budget constraint (23) determines

$$x^*(\theta) = \frac{y_1(\theta) - c_1^*(\theta)}{D} + k_0(\theta)$$

and we obtain, from the first-period budget constraint (22),

$$T_0(\theta) = y_0(\theta) + px^*(\theta) - c_0^*(\theta)$$

However, we will also consider alternative normalizations later on when this is convenient.

3.2 Taxing changing asset prices

We now consider how the tax $T_0(\theta)$ on a given investor should change in response to an asset price change. We start with the general case, which allows for asset prices driven by a mixture of discount rate and dividend changes. The two experiments discussed in Section 1.3 will then emerge as special (extremal) cases. We assume that prices and dividends are originally at some baseline levels \bar{p} and \bar{D} and then change discretely by some amounts Δp and ΔD . Importantly, at the original prices and dividends, a tax system $\bar{T}_0(\theta)$ is in place that optimally redistributes across investors and implements the corresponding Pareto efficient allocation. To obtain clean results, we will assume that preferences satisfy (6) in the following.

Proposition 1. *Let the asset price change from \bar{p} to $p = \bar{p} + \Delta p$ and dividends from \bar{D} to $D = \bar{D} + \Delta D$. Then the optimal tax $T_0(\theta)$ (when $T_1(\theta) = 0$) is given by*

$$\begin{aligned} T_0(\theta) &= \bar{T}_0(\theta) + \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D - \Omega(\theta) \left[\bar{X}\Delta p + \frac{p}{D}\bar{K}_1\Delta D \right] \\ &= \bar{T}_0(\theta) + x(\theta)\Delta p + \frac{\bar{p}}{D}k_1(\theta)\Delta D - \Omega(\theta) \left[X\Delta p + \frac{\bar{p}}{D}K_1\Delta D \right] \end{aligned}$$

where

- $\bar{T}_0(\theta)$ is the optimal tax at the initial price \bar{p} and dividends \bar{D} ,

- $\bar{x}(\theta)$ ($x(\theta)$) are investor θ 's asset sales at the initial price \bar{p} and dividends \bar{D} (new price p and dividend D) and \bar{X} (X) are the corresponding aggregate asset sales,
- $\bar{k}_1(\theta)$ ($k_1(\theta)$) are investor θ 's second-period asset holdings at the initial price \bar{p} and dividends \bar{D} (new price p and dividend D) and \bar{K}_1 (K_1) are the corresponding aggregate asset holdings,
- and $\Omega(\theta)$ is defined in (21).

Slutsky Compensation. To build intuition for this result, it is helpful to relate it to the concept of “Slutsky compensation,” which is sometimes used to define income and substitution effects of price changes. Slutsky compensation is defined as the change in the investor’s total budget (i.e., the change in initial endowment y_0) that keeps the initial consumption bundle affordable at the new prices (e.g. Mas-Colell et al., 1995, pp. 29-30).¹⁶ Using this idea, we have the following lemma:¹⁷

Lemma 1. *When the asset price changes from \bar{p} to $p = \bar{p} + \Delta p$ and dividends change from \bar{D} to $\bar{D} + \Delta D$, the corresponding Slutsky compensation $\Delta y_0(\theta)$ is given by*

$$\Delta y_0(\theta) = -\bar{x}(\theta)\Delta p - \frac{p}{D}\bar{k}_1(\theta)\Delta D.$$

This reveals that the first part of the optimal tax change characterized in the first equation in Proposition 1 coincides with the Slutsky compensation for the underlying price and dividend change. It is useful to organize the interpretation of the result and why it is related to Slutsky compensation along the special cases from Section 1.3.

3.2.1 Special Case 1: Only Discount Rate Changes

This is the first experiment discussed in Section 1.3, namely, an asset price change Δp exclusively driven by a change in the discount rate and hence holding dividends constant $\Delta D = 0$. Then we immediately obtain the following corollary of Proposition 1:

Corollary 1. *Let the asset price change from \bar{p} to $p = \bar{p} + \Delta p$ holding dividends fixed $D = \bar{D}$. Then the optimal tax $T_0(\theta)$ is given by*

$$T_0(\theta) = \bar{T}_0(\theta) + \bar{x}(\theta)\Delta p - \Omega(\theta)\bar{X}\Delta p = \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p.$$

The first equation in Corollary 1 shows that the change in the optimal tax in response to an asset price change is closely related to the Slutsky compensation from Lemma 1. Indeed, the two would coincide if there were no aggregate asset sales with the “rest of the world.” The

¹⁶The standard use of Slutsky compensation is to compute “Slutsky-compensated demand.” In particular, the difference between Slutsky-compensated demand at the new prices and demand at the old prices is one definition of the substitution effect. An alternative definition of income and substitution effects is based on “Hicksian compensation,” which is the change in total budget that restores the original level of utility. We thank Dejanir Silva for pointing out the connection of the welfare gains formula (18) to the Slutsky compensation idea. See Caramp and Silva (2023) for a related result in the context of monetary policy transmission via asset prices.

¹⁷The term “Slutsky compensation” is normally reserved for pure price changes. Here and elsewhere we use it to also refer to compensation of dividend changes ΔD .

change in $T_0(\theta)$ makes investor θ 's original consumption allocation just affordable again, and then redistributes the aggregate capital gains in the optimal way, determined by the Pareto weights $\omega(\theta)$ and the curvature of preferences γ , as summarized by the welfare weights $\Omega(\theta)$.

According to Lemma 1, to make investors' initial consumption bundle just affordable, buyers (i.e., $\bar{x}(\theta) < 0$) are compensated for the price increase (subsidized) whereas sellers (i.e., $\bar{x}(\theta) > 0$) are taxed. Figure 2 provides a graphical representation of Slutsky compensation based on the Fisher diagram, the standard graphical apparatus for intertemporal consumption choice problems. However, we include only the budget sets, and omit the indifference curves. Panel (a) plots the case of a seller while panel (b) plots that of a buyer. In both panels, the steeper solid line is the budget constraint at the initial asset price \bar{p} and the dashed line is that at the new, higher price p . A change in the asset price rotates the budget constraint through the endowment point, with an increase in price generating a shallower budget line (the slope is $-D/p$).¹⁸ A reference line is drawn through the initial consumption allocation (\bar{c}_0, \bar{c}_1) with the slope of the new budget line. The horizontal shift between the dashed line and this parallel reference line is the amount of resources needed to be added or subtracted in period 0 to afford the initial consumption allocation at the new prices. This is the Slutsky compensation. For the seller (left panel), the rise in price moves the initial consumption point into the interior of the budget set, implying a negative Slutsky compensation. The converse is true for the buyer.

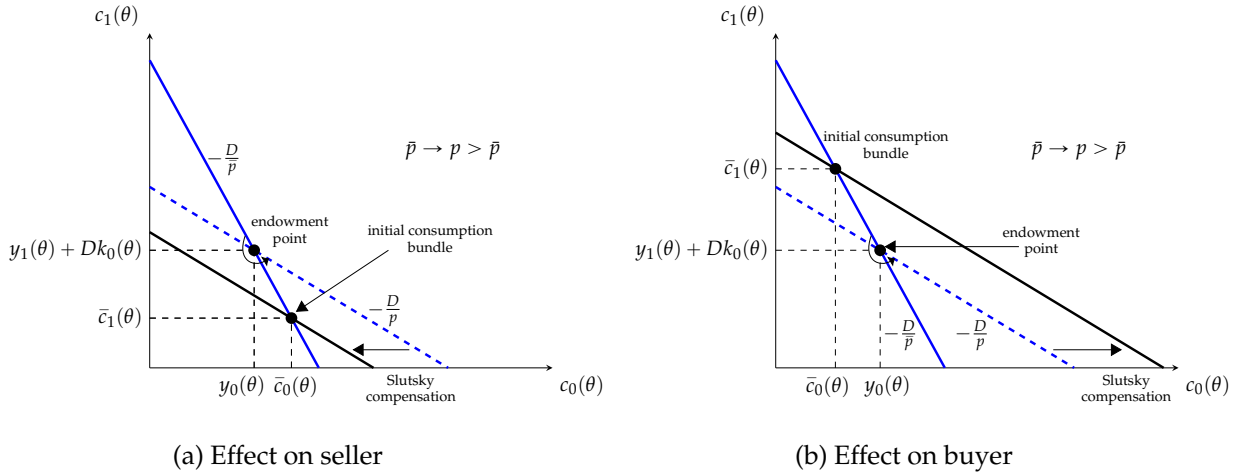


Figure 2: Slutsky compensation after a pure asset-price increase (Special Case 1)

Notes: The figure depicts Slutsky compensation in response to an asset price increase. In both panels, the solid red line is the initial budget line, with the endowment and consumption points marked. The shallower dashed line through the endowment point is the new budget line after the price change. The solid black line parallel to the dashed line is the budget line after the Slutsky compensation, which by definition contains the initial consumption allocation at the new prices. Panel (a) depicts an initial seller of the asset while Panel (b) depicts a buyer.

The intuition for why the Slutsky compensation is relevant for the optimal tax change in Corollary 1 is that a pure discount rate change does not change aggregate resources other than through trade with the rest of the world; hence, in a closed economy, the initial consumption level of every individual remains the relevant target for optimal policy, which is what precisely

¹⁸To see that the slope of the present-value budget constraint (16) is $-D/p$ and that the endowment point is given by $c_0(\theta) = y_0(\theta)$ and $c_1(\theta) = y_1(\theta) + Dk_0(\theta)$ as in Figure 2, it is useful to write it as $c_1(\theta) = \frac{D}{p}(y_0(\theta) - c_0(\theta)) + y_1(\theta) + Dk_0(\theta)$.

what Slutsky compensation is designed to deliver. The additional term in Corollary 1 then captures the optimal distribution of the aggregate gains, which is additively separable from the individual compensation given our preference assumption.

The second equation in Corollary 1 shows that the optimal tax $T_0(\theta)$ can also be written in terms of asset sales at the new asset price p . For example, if $x(\theta) > 0$ and $\Delta p > 0$, then $T_0(\theta)$ effectively taxes the *realized* capital gains of investor θ . Because of the lump-sum nature of the tax system, these gains are in fact taxed away completely, at a rate of 100%. Note that $x(\theta)$ are the new asset sales not only at the new price but also at the new taxes. In certain cases, the old and new asset sales coincide, $x(\theta) = \bar{x}(\theta)$.¹⁹

When $X = 0$, as would be the case when $\beta\bar{D}/p = 1$ and $Y_0 = Y_1 + \bar{D}K_0$ or in the closed economy we consider in Section 6.1, the tax

$$T_0(\theta) = \bar{T}_0(\theta) + x(\theta)\Delta p$$

corresponds to a realization-based capital gains tax (relative to the reference price \bar{p}), akin to the kind of capital gains taxes implemented in many countries. However, our tax formula is not limited to when the investor sells the asset ($x(\theta) > 0$) and realizes a gain ($\Delta p > 0$). It also prescribes to compensate realized capital losses ($x(\theta) > 0$ and $\Delta p < 0$) as well as purchasing gains and losses (when $x(\theta) < 0$). For instance, when the investor purchases the asset and its price falls, she benefits from a "purchasing gain" $x(\theta)\Delta p > 0$, which is also taxed away. Generally, the optimal tax base is given by what we refer to as the "trading gains and losses."

Importantly, when the investor does not trade ($x(\theta) = 0$), no tax change is triggered by the asset price change (except for a redistribution of the potential aggregate capital gains or losses $X\Delta p$). This reveals another difference from typical real-world capital gains taxes: the optimal tax conditions on *net* transactions only. For instance, if an individual sells a house and then buys another house of the same quality and price, and house prices go up, she realizes a capital gain on the sold house and a purchasing loss on the purchased house, which cancel out. By contrast, since typical tax systems, in practice, do not contain the second component (i.e., the subsidy on the purchasing loss), they would only tax the (gross) realized capital gains from the first transaction.

3.2.2 General Case

We can now return to the general case in Proposition 1. In addition to the terms discussed so far, from the pure price change Δp , new terms capturing the dividend change ΔD emerge. According to both formulas in Proposition 1, the additional dividend income, discounted back to period 0, must also be taxed away, and the aggregate dividend income change is redistributed optimally according to the welfare weights. In other words, the tax/subsidy on trading gains and losses must be complemented by a dividend income tax if dividends increase, and a subsidy if they decrease.

¹⁹In particular, this happens in the closed economy of Section 6.1 in which aggregate asset sales are zero, $X = \bar{X} = 0$. In this case, optimal policy simply takes everyone back to their initial consumption allocation, which implies $x(\theta) = \bar{x}(\theta)$.

This is again closely related to the Slutsky compensation in Lemma 1. Intuitively, investors with asset holdings $\bar{k}_1(\theta) > 0$ benefit from a higher dividend $\Delta D > 0$ and therefore need to be taxed in order to make their initial consumption bundle just affordable. Graphically, the combination of rising asset prices and rising dividends means that the budget line not only rotates around the endowment point but also shifts outwards due to the rising dividends—see Figure 3a.

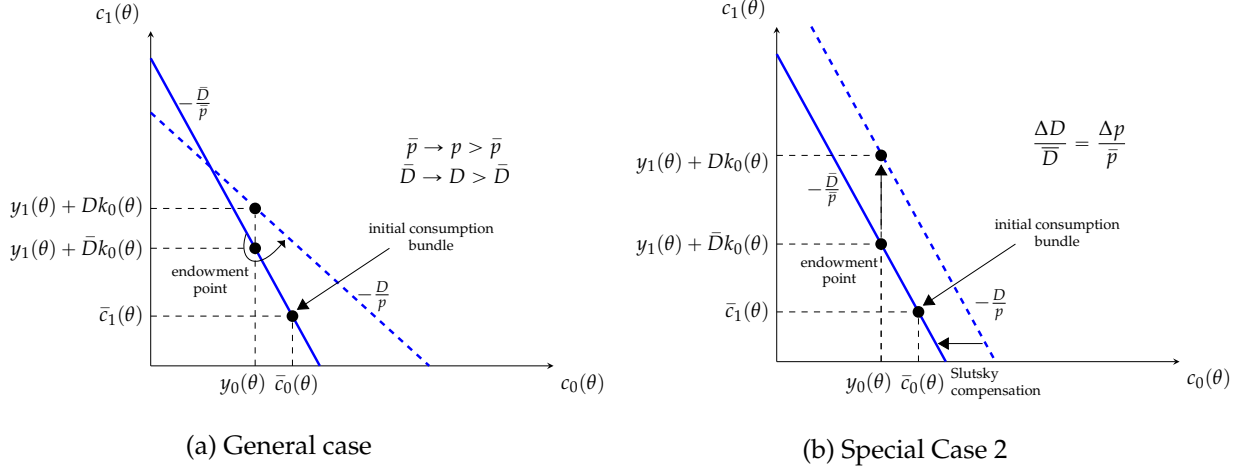


Figure 3: Slutsky compensation of combined asset-price and cash-flow changes

Notes: This figure is similar in composition to Figure 2, but allows for both price and dividend changes. Relative to the previous figure, in this case the endowment point shifts by the change in dividend income. The figure omits the new budget line absent the Slutsky compensation.

While the intuition is therefore similar to the welfare gains formula (18), an important difference is that the Slutsky compensation argument follows exclusively from investors' budget constraints at the two prices. As a result, assumptions on preferences or the optimality of the initial allocation (used in applying the envelope theorem in equation (18)) play no role. Since budget constraints are linear in prices, Lemma 1 holds for arbitrary non-infinitesimal asset price and dividend changes. This property translates to the optimal tax result in Proposition 1.

Similar to the special case in Corollary 1, Proposition 1 shows that $T_0(\theta)$ can be written both in terms of asset sales $\bar{x}(\theta)$ and asset holdings $\bar{k}_1(\theta)$ under the *old* prices (in which case dividend income must be discounted using the *new* discount rate p/D) or in terms of asset sales $x(\theta)$ and asset holdings $k_1(\theta)$ under the *new* prices (in which case the *old* discount rate \bar{p}/\bar{D} must be used). In fact, when we allow the lump-sum taxes in both periods to adjust (rather than using the normalization $T_1(\theta) = 0$), we can write the optimal tax as

$$T_0(\theta) = \bar{T}_0(\theta) + \bar{x}(\theta)\Delta p - \Omega(\theta)\bar{X}\Delta p = \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p$$

$$T_1(\theta) = \bar{T}_1(\theta) + \bar{k}_1(\theta)\Delta D - \Omega(\theta)\bar{K}_1\Delta D = \bar{T}_1(\theta) + k_1(\theta)\Delta D - \Omega(\theta)K_1\Delta D,$$

where $\{\bar{T}_t(\theta)\}$, $t = 0, 1$, are some taxes that implement the optimum at the old prices and dividends. Hence $T_0(\theta)$ deals with the pure asset price change in the form of a realization-based tax on capital gains just like in Corollary 1, whereas $T_1(\theta)$ acts as a tax on the changed dividend

income in $t = 1$. In particular, no discounting is needed under this alternative normalization.²⁰

3.2.3 Special Case 2: Only Cash Flow Changes

We are now ready to consider the second extreme case from Section 1.3 where asset prices change exclusively because of a change in future dividends. For simplicity, we return to the normalization $T_1(\theta) = 0$.

Corollary 2. *Let the asset price change from \bar{p} to $p = \bar{p} + \Delta p$ and let dividends change from \bar{D} to $D = \bar{D} + \Delta D$ such that $\Delta D / \Delta p = \bar{D} / \bar{p}$. Then the optimal tax $T_0(\theta)$ is given by*

$$T_0(\theta) = \bar{T}_0(\theta) + k_0(\theta)\Delta p - \Omega(\theta)K_0\Delta p$$

Since dividends and the asset price grow by the same percentage, the asset return $R_1 = D/p$ remains unchanged, i.e., the new return $R_1 = (\bar{D} + \Delta D) / (\bar{p} + \Delta p)$ equals the old return $\bar{R}_1 = \bar{D} / \bar{p}$. According to the corollary, the optimal tax $T_0(\theta)$ then taxes the investor's individual wealth gains $k_0(\theta)\Delta p$ due to the asset price change. Hence, in this case, the first-best optimal tax system conditions on the investor's *unrealized* capital gains. This tax base is therefore consistent with an accrual-based capital gains tax, as under the Haig-Simons comprehensive income tax (Haig, 1921, and Simons, 1938), or a wealth tax.

Of course, since Proposition 1 continues to apply, we could still express the optimum as a combination of a tax on trading gains and a dividend income tax. Why do the two collapse to the accrual-based tax in Corollary 2, which only depends on initial wealth in $t = 0$? The reason can be understood as follows: If the investor sells all her assets, then $x(\theta) = k_0(\theta)$ and $k_1(\theta) = 0$. Hence, there is no dividend income in $t = 1$, and realized capital gains in $t = 0$ are given by $k_0(\theta)\Delta p$, just as in the corollary. Now suppose instead that the investor decides not to sell all her assets. This results in some dividend income in $t = 1$ (since now $k_1(\theta) = k_0(\theta) - x(\theta) > 0$), but at the same time in reduced realized capital gains in $t = 0$. When the price and dividend changes happen to be proportional, the two effects exactly offset each other and the overall income change is still given by $k_0(\theta)\Delta p$, no matter how much the individual sells. Formally, we can use the first equation in Proposition 1, for instance, to obtain

$$\bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D = \bar{x}(\theta)\Delta p + \frac{p}{D}(k_0(\theta) - \bar{x}(\theta))\Delta p \frac{D}{p} = k_0(\theta)\Delta p$$

Since this holds for all investors, the aggregate quantities collapse in the same way, so Corollary 2 obtains.

²⁰The same is true for the Slutsky compensation. When we allow compensation in both periods, i.e. $\Delta y_1(\theta)$ in addition to $\Delta y_0(\theta)$, we can write

$$\Delta y_0(\theta) = -\bar{x}(\theta)\Delta p \quad \text{and} \quad \Delta y_1(\theta) = -\bar{k}_1(\theta)\Delta D.$$

In general, only the present value of $\Delta y_0(\theta)$ and $\Delta y_1(\theta)$ is pinned down. In particular, one can show that any combination satisfying

$$\Delta y_0(\theta) + \frac{p}{D}\Delta y_1(\theta) = -\bar{x}(\theta)\Delta p - \frac{p}{D}\bar{k}_1(\theta)\Delta D$$

achieves the Slutsky compensation goal of making the initial consumption bundle just affordable.

Figure 3b relates this case graphically to the corresponding Slutsky compensation. In contrast to Figures 2 and 3a, the budget line does not change slope (which remains unchanged at $-\bar{D}/\bar{p}$) and instead shifts outward. Specifically, the increase in dividends means that the endowment point $(y_0(\theta), y_1(\theta) + Dk_0(\theta))$ simply shifts upward. In the lifetime budget constraint (16), the return D/p remains unchanged and therefore the only effect of the joint asset price and dividend change is the revaluation of initial wealth $pk_0(\theta)$.

While this special case, where the asset price change is accompanied by a dividend change such that the discount rate remains fixed, therefore provides a justification for using the Haig-Simons income concept as the tax base, this logic demonstrates that it is knife-edge. Whenever capital gains are not entirely due to dividend changes, i.e. as soon as discount rate changes are part of the story as well—in the Fisher diagram of Figures 2 and 3, as soon as the budget line rotates even a little bit—the additional dividend income and capital gains no longer cancel out. Moreover, we will show in Section 6.2 that the cancellation result will break down, even when asset prices are exclusively driven by dividend changes, in a richer model with heterogeneous returns.

3.3 Discussion

Baseline asset price. Proposition 1 characterizes the change in optimal taxes in response to an asset price and dividend change. Thus, it captures a comparative-static comparison: how taxes should differ in an economy with old prices and dividends \bar{p} and \bar{D} compared to another economy with new prices and dividends p and D . This also affects the interpretation of the “trading gains and losses” $x(\theta)\Delta p$, for instance. While they bear similarities to realized capital gains (in case of an asset sale), an important difference is that computing the price change Δp requires fixing some baseline price \bar{p} . Notably, as will become even clearer in the multi-period extension in Section 5, this baseline price does not necessarily coincide with the historical basis at which the investor purchased the asset. Instead, one needs to decide which price (and dividend) change the tax system should compensate. This becomes particularly clear in the case of a purchasing gain or loss (with $x < 0$): in this case, Proposition 1 prescribes a tax or subsidy, but since the investor has not owned the asset prior to purchasing it, there is no historical basis to go back to when computing the price change.

One example for a baseline price and dividend would be to imagine that there is some ex-ante uncertainty about them, and \bar{p} and \bar{D} represent the corresponding means. In this case, the old allocation can be interpreted as the optimum under these expected prices and dividends. By contrast, the new prices and dividends p and D would be the ex-post realized ones, so the tax system is tasked to compensate the winners and losers relative to the ex-ante expectations. Another natural benchmark would be the hypothetical asset price path if discount rates had not fallen, but we also discuss other examples in Section 5.

Finally, when comparing to the models of capital taxation with linear savings technology, Special Case 2 is the relevant baseline (since it holds the rate of return fixed). The same is true for the literature based on the neoclassical growth model along a balanced growth path.

Baseline taxes. The comparison in Proposition 1 also assumes that, both under the old and new prices and dividends, taxes are set optimally for given Pareto weights. Hence, if taxes under the old prices and dividends were not set optimally (or based on different Pareto weights), we can always decompose the overall change in taxes into two steps: First, holding old prices and dividends fixed, a reform of the old taxes towards the optimum according to the new Pareto weights. Second, holding Pareto weights fixed, a move from the optimum under the old prices and dividends towards the optimum under the new ones. Our analysis isolates the second step. The first step has nothing to do with asset prices and is completely standard, namely, a tax reform moving the allocation from the interior of the Pareto frontier (or along the frontier) towards a particular point on that frontier in a given economy.

Endogenous payout policy and share repurchases. Businesses have control over their dividend payments and may have alternative means of distributing their profits to shareholders, specifically via share repurchases. Appendix B.4 therefore provides an alternative, capital-structure neutral formulation of our setup in which such distinctions are immaterial. This is particularly important when we study the second-best problem with distortive taxation later on because such a neutral formulation avoids creating incentives for firms to distort their capital structure. The key idea of this formulation is to consolidate the firm and investor budget constraints, in particular to consider profits net of investment as the relevant measure of cash flows D_t regardless of whether they are distributed to investors via dividend payments or share repurchases and to consider the firm's total value as the relevant measure of the share price p_t .

Owner-occupied housing. Owner-occupied housing generates a utility flow of housing services and implementing our tax formula therefore requires valuing this "utility dividend." The appropriate solution is to use a rental equivalence approach that measures the dividend D as imputed rents, i.e. to value owner-occupied housing services as the rental income the homeowner could have received if the house had been let out. Thus, if part of the house-price increase in New York City was due to the city's amenities improving, rents would rise so that $\Delta D_t > 0$ in addition to $\Delta p_t > 0$ and our formula would prescribe taxing the additional imputed rents. This approach is already used by some countries including Switzerland ("Eigenmietwert") and Denmark.

3.4 Alternative implementations: taxes on total capital income or expenditure

In Proposition 1, we have expressed the first-best tax response to asset price and dividend changes in terms of investors' realized capital gains and additional dividend income. We now show that the optimal tax change can be equivalently understood in two alternative ways: one based on total capital income and another one based on consumption.

3.4.1 A tax on total capital income.

In Section 1, we demonstrated that equation (5) is an equivalent way of writing the budget constraint (2) when using an investor's market value of wealth $a_0(\theta) \equiv p_{-1}k_0(\theta)$ in period 0 and $a_1(\theta) \equiv pk_1(\theta)$ in period 1. This implies that an alternative way of writing the first-best tax response in Proposition 1 is in terms of these wealth holdings and the changes in the total returns R_0 and R_1 , as the following proposition shows.

Proposition 2. *Let the asset price change from \bar{p} to $p = \bar{p} + \Delta p$ and dividends from \bar{D} to $D = \bar{D} + \Delta D$ resulting in return changes $\Delta R_0 = R_0 - \bar{R}_0$ and $\Delta R_1 = R_1 - \bar{R}_1$. Then the optimal tax $T_0(\theta)$ (when $T_1(\theta) = 0$) is given by*

$$\begin{aligned} T_0(\theta) &= \bar{T}_0(\theta) + a_0(\theta)\Delta R_0 + \frac{1}{R_1}\bar{a}_1(\theta)\Delta R_1 - \Omega(\theta) \left[A_0\Delta R_0 + \frac{1}{R_1}\bar{A}_1\Delta R_1 \right] \\ &= \bar{T}_0(\theta) + a_0(\theta)\Delta R_0 + \frac{1}{R_1}a_1(\theta)\Delta R_1 - \Omega(\theta) \left[A_0\Delta R_0 + \frac{1}{R_1}A_1\Delta R_1 \right] \end{aligned}$$

where $a_1(\theta)$ ($\bar{a}_1(\theta)$) is investor θ 's period-1 wealth at the new (old) returns ($a_0(\theta) = \bar{a}_0(\theta)$ since p_{-1} is fixed) and A_0 , A_1 and \bar{A}_1 are the corresponding aggregate wealth holdings.

At first glance, the tax change in Proposition 2 appears related to a Haig-Simons notion of income: in each period, the additional total capital income $a_t(\theta)\Delta R_t$, including unrealized gains, is taxed (and, just like in Proposition 1, this can be written both in terms of the old or new wealth holdings and the aggregate income is redistributed optimally according to the weights $\Omega(\theta)$). However, there is an important difference. It is easiest to see this by considering an increase in the asset price p holding dividends fixed (Special Case 1). In this case, we have

$$\Delta R_0 = \frac{\Delta p}{p_{-1}} > 0 \quad \text{and} \quad \Delta R_1 = \frac{D}{p} - \frac{D}{\bar{p}} < 0$$

since $p = \bar{p} + \Delta p > \bar{p}$. Hence, the investor faces a tax in period 0 (due to the higher return from the increased asset price) but a subsidy in period 1. The reason for the latter is that, whereas the asset price has increased, dividends have not, so the dividend-price ratio and thus the return in period 1 has been reduced, which needs to be compensated.

While the former tax increase (due to the unrealized capital gains in the initial period) indeed corresponds to a Haig-Simons income tax, the latter subsidy (due to the lower dividend-price ratio subsequently) does not. But Proposition 2 shows that they belong together. Therefore, due to these opposing effects, the total change in taxes is generally ambiguous. In fact, we know from Proposition 1 that it depends solely on whether the investor is a buyer or seller. For instance, when $x(\theta) = 0$ so the individual is not trading, the additional tax in period 0 and the subsidy in period 1 precisely cancel out.²¹

The intuition for Proposition 2 can again be linked to the Slutsky compensation. In other

²¹In the multi-period model, a one-off, permanent increase in the asset price would trigger a one-off tax followed by a subsidy forever, in a way that their present value sum is zero for an investor who is not trading (see Proposition 9 in Appendix D.3). Thus, the alternative implementation in Proposition 2 can lead to very volatile taxes over time compared to Proposition 1.

words, it is possible to re-write the Slutsky compensation in Lemma 1 based on the investor's wealth holdings and total return changes, as follows:

Lemma 2. *When the asset price changes from \bar{p} to $p = \bar{p} + \Delta p$ and the dividend changes from \bar{D} to $\bar{D} + \Delta D$, the corresponding Slutsky compensation $\Delta y_0(\theta)$ is given by*

$$\Delta y_0(\theta) = -a_0(\theta)\Delta R_0 - \frac{1}{R_1}\bar{a}_1(\theta)\Delta R_1.$$

3.4.2 An expenditure tax

There is a long-standing debate in public finance about the potential advantages of taxes on consumption or expenditures, notably in the context of capital gains. For instance, when discussing the Haig-Simons income concept (which is typically understood to include consumption plus unrealized capital gains), Kaldor (1955, p. 44) writes:

“We may now turn to the other type of capital appreciation which reflects a fall in interest rates rather than the expectation of higher earning power. [...] The rise in capital values in this case [comes] without a corresponding increase in the flow of real income accruing from that wealth. [...] For in so far as a capital gain is realized [...] the benefit derived from the gain is equivalent to that of any other casual profit. If however it is not so realized, there is clearly only a smaller benefit. [Therefore] treating the two kinds of capital gains in the same way is not an equitable method of measuring taxable capacities.”

Given this problem with using Haig-Simons income as the tax base, Kaldor instead advocates for an expenditure-based tax. This raises the question whether the optimal tax response to changing asset prices and dividends in Proposition 1 could also be understood as a consumption-based (or, equivalently, an expenditure-based) tax, independent of the source of capital gains. We formalize this conjecture in the following proposition. To do so, denote by $\{\bar{c}_t(\theta)\}$, $t = 0, 1$, the optimal consumption allocation at the old prices and dividends \bar{p} and \bar{D} (i.e., the solution to the Pareto problem (19)), and by $\{c_t(\theta)\}$ at the new prices and dividends $p = \bar{p} + \Delta p$ and $D = \bar{D} + \Delta D$. Let $\{\bar{T}_t(\theta)\}$, $t = 0, 1$, be some lump-sum taxes that implement the optimum at the old prices and dividends. Finally, let $\hat{c}_t(\theta)$, $t = 0, 1$, denote investor θ 's *individually* optimal consumption allocation under the new prices and dividends but when taxes are held fixed at the old level. Formally, $\hat{c}_t(\theta)$ solves

$$\max_{c_0(\theta), c_1(\theta), x(\theta)} U(c_0(\theta), c_1(\theta)) \quad \text{s.t.} \quad (22) \text{ and } (23)$$

when taxes are given by $\{\bar{T}_t(\theta)\}$. Then we have the following result:

Proposition 3. *Suppose asset prices increase from \bar{p} to $p = \bar{p} + \Delta p$ and dividends from \bar{D} to $D = \bar{D} + \Delta D$. The optimal consumption allocation at the new prices and dividends $\{c_t(\theta)\}$, $t = 0, 1$, can be implemented with taxes given by*

$$T_t(\theta) = \bar{T}_t(\theta) + \Delta \hat{c}_t(\theta) - \Omega(\theta)\Delta C_t, \quad t = 0, 1$$

where $\Delta\hat{c}_t(\theta) \equiv \hat{c}_t(\theta) - \bar{c}_t(\theta)$ and

$$\Delta C_t = \int c_t(\theta)dF(\theta) - \int \bar{c}_t(\theta)dF(\theta).$$

Hence, the new optimum can be implemented with a combination of a lump-sum tax equal to $\Delta\hat{c}_t(\theta)$, which is the amount by which individuals would have changed their consumption after the price and dividend change if taxes had stayed at their old level $\bar{T}_t(\theta)$, and a transfer equal to the difference between the old and new desired consumption $\Delta c_t(\theta) = \Omega(\theta)\Delta C_t$. Notably, if the parameter changes Δp and ΔD are "zero-sum," so that optimal aggregate consumption C_t does not change, then a pure expenditure-based tax is sufficient. In line with Kaldor's logic, just like Proposition 1, this works for any combination of asset price and dividend changes. Indeed, we show in Appendix B that

$$\Delta\hat{c}_0(\theta) + \frac{p}{D}\Delta\hat{c}_1(\theta) = \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D, \quad (24)$$

so the present value of the consumption change (holding taxes fixed) is directly linked to the capital gains and change in dividend income when p and D change. For instance, the investors who would increase their consumption in response to a pure asset price increase (in the absence of a further tax change) are precisely those who sell the asset, and vice versa.

4 Second best

We now turn to the case where the government's tax instruments are more limited. Specifically, rather than having access to lump-sum taxes, it is restricted to taxes that condition on investors' choices, such as their asset sales, wealth, or consumption. Hence, these taxes will distort individuals' behavior, inducing the classic tradeoff between redistribution and efficiency. While this will lead to optimal marginal tax rates less than 100%, our main conclusion will be that the previous results on the tax base generalize in a natural way.

4.1 Mirrlees Problem

An asset sales tax. We begin with the case of a (non-linear) tax $T_x(px)$ on asset sales, paid in period 0, which is similar to a realization-based capital gains tax. Hence, the investors' budget constraints become

$$\begin{aligned} c_0(\theta) &= y_0(\theta) + px(\theta) - T_x(px(\theta)) \\ c_1(\theta) &= y_1(\theta) + D(k_0(\theta) - x(\theta)). \end{aligned}$$

In the language of Mirrlees (1971), this corresponds to a situation where $x(\theta)$ (and hence $z_x(\theta) \equiv px(\theta) - T_x(px(\theta))$) is observable but $k_0(\theta)$, $y_0(\theta)$ and $y_1(\theta)$ are not. The incentive

constraints are therefore

$$\begin{aligned} U(\theta) &\equiv U(z_x(\theta) + y_0(\theta), D(k_0(\theta) - x(\theta)) + y_1(\theta)) \\ &\geq U(z_x(\hat{\theta}) + y_0(\theta), D(k_0(\theta) - x(\hat{\theta})) + y_1(\theta)) \quad \forall \theta, \hat{\theta}. \end{aligned} \quad (25)$$

Abstracting from bunching, we work with the local version of the incentive constraints. By the envelope theorem,

$$U'(\theta) = U_{c_0}(c_0(\theta), c_1(\theta))y'_0(\theta) + U_{c_1}(c_0(\theta), c_1(\theta))(Dk'_0(\theta) + y'_1(\theta)) \quad \forall \theta. \quad (26)$$

Hence, the second-best problem corresponding to the optimal asset sales tax is

$$\max_{\{c_0(\theta), c_1(\theta), U(\theta)\}} \int \omega(\theta) U(\theta) dF(\theta) \quad (27)$$

subject to the resource constraint (17), the incentive constraints (26), and $U(\theta) = U(c_0(\theta), c_1(\theta))$. Given a constrained optimal allocation, we can then back out the optimal marginal asset sales tax $T'_x(px(\theta))$ from the individual Euler equations:

$$T'_x(px(\theta)) = 1 - \frac{D U_{c_1}(c_0(\theta), c_1(\theta))}{p U_{c_0}(c_0(\theta), c_1(\theta))}.$$

A wealth tax. Alternatively, consider a tax $T_k(pk_1(\theta))$ on investors' wealth in period 1.²² This corresponds to a setting where $k_1(\theta)$ (and hence $z_k(\theta) \equiv Dk_1(\theta) - T_k(pk_1(\theta))$) is observable, resulting in the global incentive constraints

$$\begin{aligned} U(\theta) &\equiv U(p(k_0(\theta) - k_1(\theta)) + y_0(\theta), z_k(\theta) + y_1(\theta)) \\ &\geq U(p(k_0(\theta) - k_1(\hat{\theta})) + y_0(\theta), z_k(\hat{\theta}) + y_1(\theta)) \quad \forall \theta, \hat{\theta}. \end{aligned} \quad (28)$$

The local incentive constraints can therefore be written in the same general form as in the case of the asset sales tax, namely

$$U'(\theta) = U_{c_0}(c_0(\theta), c_1(\theta))A(\theta) + U_{c_1}(c_0(\theta), c_1(\theta))B(\theta) \quad \forall \theta, \quad (29)$$

with the only difference that, now,

$$A(\theta) = pk'_0(\theta) + y'_0(\theta) \quad \text{and} \quad B(\theta) = y'_1(\theta).$$

Hence, the second-best problem for the wealth tax is the same as above, except for the incentive constraints (29) instead of (26).²³ For a given optimal allocation, we can again compute the

²²This is equivalent to a tax on dividend income $Dk_1(\theta)$ since dividends D are the same for all investors in our baseline model. By contrast, a tax on period-0 wealth $pk_0(\theta)$ would be lump-sum and return us to the first-best case when $k_0(\theta)$ is invertible.

²³In fact, as can be seen immediately from the incentive constraints, the two second-best problems are identical when investors only differ in their incomes $y_0(\theta)$ and $y_1(\theta)$ but not in their initial wealth $k_0(\theta)$. In other words, in this case, a wealth tax and an asset sales tax are two decentralizations of the same optimal allocation.

optimal marginal wealth tax based on the Euler equation:

$$T'_k(pk_1(\theta)) = \frac{D}{p} - \frac{U_{c_0}(c_0(\theta), c_1(\theta))}{U_{c_1}(c_0(\theta), c_1(\theta))}.$$

Other taxes. We can allow for other taxes instruments that condition on investors' choices. For example, the second-best problems for consumption taxes are very similar. Notably, we show in Appendix C that the general incentive constraints (29) still apply, with the only difference that the terms $A(\theta)$ and $B(\theta)$ are modified. In the case of a nonlinear tax on c_0 , we have $A(\theta) = 0$ and

$$B(\theta) = Dk'_0(\theta) + \frac{D}{p}y'_0(\theta) + y'_1(\theta),$$

whereas a tax on c_1 implies

$$A(\theta) = pk'_0(\theta) + y'_0(\theta) + \frac{p}{D}y'_1(\theta)$$

and $B(\theta) = 0$. More generally, this structure extends to any such second-best problem, including when combinations of taxes are available (such as an asset sales tax in period 0 plus a wealth tax in period 1).

Solution. In Appendix C, we show how to find the Mirrleesian Pareto optima for all these tax instruments. We can then use this solution to show how the optimal tax schedule changes in response to an asset price change. In the next subsection, we numerically compute these changes for the CES preferences defined in equation (6).

4.2 Taxing changing asset prices

We consider an example economy with investors θ uniformly distributed on the unit interval and $y_0(\theta) = 1 - \theta$, $y_1(\theta) = \theta$ and $k_0(\theta) = 0.1$ for all $\theta \in [0, 1]$. Thus, similar to Figure 1, higher- θ investors feature a more backloaded income profile (while there is no heterogeneity in the initial asset endowment), making them natural sellers (borrowers), whereas lower- θ investors are buyers (savers). As a baseline, we set the intertemporal elasticity of substitution $\sigma = 0.5$, $\gamma = 1$ (so $G(C) = \log(C)$), $\beta = 0.5$, $\bar{p} = 1$ and $\bar{D} = 2$. We compute the utilitarian optimum (with $\omega(\theta) = 1$ for all θ) for this baseline and then compare it to a situation where the asset price rises by 30% to $p = 1.3$, holding dividends \bar{D} fixed. As a result, this exercise illustrates Special Case 1, where the asset price change is entirely driven by a change in the discount rate.

Assets sales tax. The left panel in Figure 4 shows the optimal asset sales tax schedules $T_x(px)$ in both of these economies, which are decreasing. The reason is that, in this specification, higher- θ individuals have the lower present-value of income, so the direction of redistribution runs from low- to high- θ types. As discussed above, asset sales x are naturally increasing in θ , so the optimal tax schedule puts a tax on the (richer) buyers and a subsidy on the (poorer) sellers.

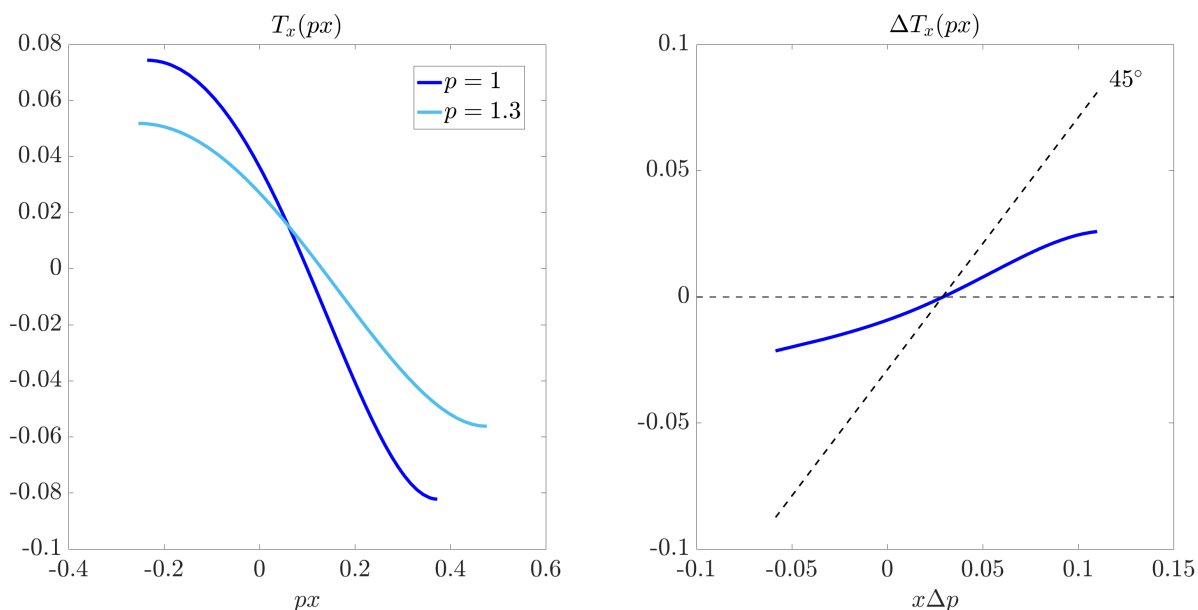


Figure 4: Optimal asset sales tax with increasing asset prices

Notes: The left panel depicts the second-best tax schedule, as a function of asset sales $px(\theta)$, for two alternative prices of the asset. The right panel depicts the difference between the schedule associated with $p=1.3$ and the baseline $p=1.0$ schedule, ΔT_x . The right panel plots $\Delta T_x(px(\theta))$ against the net trading gains $x(\theta)\Delta p$, depicting that the change in taxes increases in net gains.

Our main interest is in how the optimal tax *changes* in response to the asset price increase. This is depicted in the right panel of Figure 4, where we plot the change in the tax $\Delta T_x(px)$ as a function of the trading gains and losses $x\Delta p$. It reveals a positive relationship, just like in Corollary 1, albeit with a slope of less than one. This is intuitive: the solution now balances the optimal redistribution, which works in the same way as in the first-best case (namely, increasing the tax burden on the sellers, who gain from the asset price increase, and lowering it for the buyers), with the distortive effects of a positive marginal tax rate on investors' savings behavior. In other words, the rising asset price already accomplishes some redistribution from richer to poorer individuals in this economy, and the optimal tax policy response therefore involves less redistribution than before.

Wealth tax. Figure 5 shows the corresponding graphs for the alternative implementation of the optimum with a wealth tax. In this example, the wealth tax is increasing in period-1 wealth pk_1 : since richer, low- θ investors have a more front-loaded income stream, they will buy more assets and hence exhibit more wealth at the beginning of the second period. Thus, in terms of levels, a progressive wealth tax with a positive tax burden on wealthy individuals and a subsidy on those with low wealth (or even debt) is optimal. However, the right panel shows that the optimal response to increasing asset prices is to make the wealth tax *less* progressive. This is because, again, wealthy individuals are buyers in this case, who lose from the asset price increase, so their tax burden should fall as a compensation. Conversely, low-wealth borrowers are sellers of the asset, and hence benefit from the asset price increase, so their tax burden should increase.




Figure 5: Optimal wealth tax with increasing asset prices

Notes: The left panel depicts the second-best tax schedule, as a function of period one wealth, pk_1 , for two alternative prices of the asset. The right panel depicts the difference between the schedule associated with $p=1.3$ and the baseline $p=1.0$ schedule, ΔT_k . The right panel plots $\Delta T_k(pk_1)$ against period one wealth, depicting that the change in taxes decreases in period one wealth.

An example of such a configuration would be housing markets where relatively well-off households, who already own a house, want to upsize (for instance because of a growing family). Thus, despite being in the upper percentiles of the wealth distribution, these households are net buyers. As a result, when house prices rise, they are worse off, and introducing a progressive wealth tax in this situation would not achieve the desired direction of redistribution.

Thus, while a wealth tax in this example can be used to implement the constrained optimum, its comparative statics in response to an asset price increase may be counter-intuitive. This is because the wealth tax (or related accrual-based tax instruments) is an indirect way of targeting buyers versus sellers, which is what ultimately drives the welfare effects. By contrast, the comparative statics of the asset sales tax are always the same (as in the right panel of Figure 4) since conditioning on realized capital gains directly targets the correct tax base.

4.3 Role of the intertemporal elasticity of substitution

In Figure 6, we return to the asset sales tax and show its optimal response to an asset price increase for different values of the intertemporal elasticity of substitution (the dark blue schedule is the same as in the right panel of Figure 4). It illustrates that the optimal second-best policy converges to the first-best solution in Corollary 1, with a 100% marginal tax rate on realized capital gains, as σ approaches zero. The intuition is simply that a vanishing substitution elasticity implies a vanishing savings distortion from the tax, which therefore becomes equivalent in the limit to a lump-sum tax instrument. This demonstrates that our first-best results from the previous Section 3 are not knife-edge, but extend qualitatively to the case of more realistic and limited tax instruments as long as the distortive effects remain small. The next proposition formalizes this result.




Figure 6: Capital gains tax with a decreasing intertemporal elasticity of substitution

Denote by $\Gamma^* \equiv \{(c_0^*(\theta), c_1^*(\theta))\}$ the optimal first-best allocation solving (19) subject to (17) when preferences are given by (6) with $\sigma = 0$, so that $C(c_0, c_1) = \min\{c_0, c_1\}$. Let $V^*(\theta) \equiv C(c_0^*(\theta), c_1^*(\theta))$ for all θ . For any σ , denote by $\Gamma^M(\sigma) \equiv \{(c_0^M(\theta, \sigma), c_1^M(\theta, \sigma))\}$ the solution to the general Mirrlees problem (27) subject to (17) and (29) for the same Pareto weights $\{\omega(\theta)\}$. Finally, let $a(\theta) \equiv \min\{A(\theta), B(\theta)\}$ and $b(\theta) = \max\{A(\theta), B(\theta)\}$ for each $\theta \in [\underline{\theta}, \bar{\theta}]$. Then we obtain the following result:

Proposition 4. *If $V^*(\theta)$ satisfies $V^{*'}(\theta) \in (a(\theta), b(\theta))$ for all θ , then*

$$\lim_{\sigma \rightarrow 0} \Gamma^M(\sigma) = \Gamma^*,$$

i.e., in the Leontief limit $\sigma \rightarrow 0$, the solution to the Mirrlees problem converges to the first-best allocation.

4.4 Portfolio choice with distortive taxes and the lock-in effect

Given that we work with a single-asset environment, our analysis in this section does not allow for the possibility of taxes distorting portfolio choice. This is an important limitation that should be addressed in subsequent work by extending our analysis to multiple assets. An important case of taxes distorting portfolio choice is the "lock-in" effect emphasized in the capital gains taxation literature (see e.g. [Auerbach, 1991](#); [Constantinides, 1983](#); [Chari et al., 2005](#)). Specifically, the idea is that a realization-based tax system creates incentives to defer the liquidation of appreciated assets due to the resulting interest advantage. While we briefly revisit this issue in our analysis of the multi-period model in the next Section, a full analysis of optimal distortive taxation in the presence of the lock-in effect is left for future work.




Figure 7: Example time paths for asset price, dividends, and discount rate

5 Optimal taxation in the multi-period model

In this section, we return to the case of first-best tax instruments and show how our results on optimal taxation from the two-period model extend to the general multi-period model introduced in Section 1.

5.1 Comparing asset price and dividend paths

In its most general form, Proposition 5 below will consider arbitrary baseline time paths $\{\bar{p}_t\}_{t=0}^T$ and $\{\bar{D}_t\}_{t=0}^T$ which then change to

$$\{p_t\}_{t=0}^T = \{\bar{p}_t + \Delta p_t\}_{t=0}^T \quad \text{and} \quad \{D_t\}_{t=0}^T = \{\bar{D}_t + \Delta D_t\}_{t=0}^T.$$

To fix ideas, it is useful to consider a particular thought experiment. The economy is initially in a steady state with constant dividend \bar{D} , asset price \bar{p} and associated asset return $\bar{R} = 1 + \bar{D}/\bar{p}$. In this initial steady state, there is a tax system $\bar{T}_t(\theta)$ in place that optimally redistributes across investors, thereby implementing the Pareto optimal allocation. At time $t = 0$, the time paths of dividends and asset prices change to

$$\{p_t\}_{t=0}^T = \{\bar{p} + \Delta p_t\}_{t=0}^T \quad \text{and} \quad \{D_t\}_{t=0}^T = \{\bar{D} + \Delta D_t\}_{t=0}^T.$$

Equivalently, as in Section 1.3, we can think of the change in asset prices as being driven by changes in discount rates $\{\Delta R_{t+1}\}_{t=0}^T$ and cash flows $\{\Delta D_t\}_{t=0}^T$.²⁴ The question we consider is: how should optimal redistributive taxes respond to these capital gains?

Figure 7 illustrates one example of this thought experiment: at time zero, asset prices p_t start increasing but without any corresponding increase in cash flows D_t (panels (a) and (b)); equivalently, the asset discount rate R_t jumps up initially and then declines secularly to a lower

²⁴As discussed in Section 1.3, in equilibrium models, it is typically the time path of asset returns $\{R_{t+1}\}$ that is pinned down which, in turn, determines the asset price $p_t = \sum_{s=1}^{T-t} R_{t \rightarrow t+s}^{-1} D_{t+s}$.

long-run level (panel (c)) and this is what generates the capital gains. This example therefore corresponds to Special Case 1 discussed in Section 1.3 in which the asset price changes are driven entirely by discount rate changes.

Instead of being in a steady state initially, the economy could alternatively be on a balanced growth path (BGP) where dividends grow at some constant rate G , so $\bar{D}_t = G^t \bar{D}_0$, and rates of return are constant: $\bar{R}_{t+1} = \bar{R}$ for all t . For example, this would happen in an equilibrium model with a neoclassical production side and productivity growth of the type discussed in Section 1.4. This implies that prices also grow at a constant rate $\bar{p}_t = G^t \bar{p}_0$, consistent with a BGP.²⁵ Moreover, in this initial BGP, there is a (time-varying) tax system $\bar{T}_t(\theta)$ in place that optimally redistributes across investors. We will return to this case below.

5.2 Taxing changing asset prices and dividends

The Pareto problem with lump-sum taxes is

$$\max_{\{c_t(\theta), k_{t+1}(\theta)\}_{t=0}^T} \int \omega(\theta) U(c_0(\theta), \dots, c_T(\theta)) dF(\theta)$$

subject to (7) and the given time paths for the asset price $\{p_t\}_{t=0}^T$ and dividend $\{D_t\}_{t=0}^T$. This yields the following generalization of Proposition 1:

Proposition 5. *Suppose asset prices change by $\{\Delta p_t\}_{t=0}^T$ and dividends by $\{\Delta D_t\}_{t=0}^T$. Then optimal taxes $\{T_t(\theta)\}_{t=0}^T$ are such that*

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [\bar{T}_t(\theta) + x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Omega(\theta) (X_t \Delta p_t + K_t \Delta D_t)]$$

where $x_t(\theta) = k_t(\theta) - k_{t+1}(\theta)$.

In the multi-period model, our results from the two-period case thus apply to the present value of taxes.²⁶ So far, we have not normalized taxes, so Proposition 5 leaves room for many different implementations. A natural example is to set

$$T_t(\theta) = \bar{T}_t(\theta) + x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Omega(\theta) (X_t \Delta p_t + K_t \Delta D_t) \quad \text{for all } t, \quad (30)$$

meaning that we tax or subsidize the trading gains and losses as well as the change in dividend income period by period.

²⁵Specifically, when $\bar{D}_t = G^t \bar{D}_0$ and $\bar{R}_{t+1} = \bar{R}$ for all t , then $\bar{p}_t = \frac{G}{G-\bar{R}} \bar{D}_t$ for all t .

²⁶Proposition 5 expresses the change in the present value of taxes using sales $x_t(\theta)$ and asset holdings $k_t(\theta)$ under the new prices and dividends, but the rates of return $\bar{R}_{0 \rightarrow t}$ under the old prices and dividends. Analogously to Proposition 1, it is also possible to write the change in taxes in the opposite way, namely

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [\bar{T}_t(\theta) + \bar{x}_t(\theta) \Delta p_t + \bar{k}_t(\theta) \Delta D_t - \Omega(\theta) (\bar{X}_t \Delta p_t + \bar{K}_t \Delta D_t)]$$

Special Case 1. It is instructive to consider the first extreme case from Section 1 where asset prices change exclusively because of a change in discount rates, i.e. $\Delta D_t = 0$ for all t . For example, as in Figure 7, the economy could initially be in a steady state with constant dividend \bar{D} , asset price \bar{p} and discount rate \bar{R} and with a tax system $\bar{T}_t(\theta)$ that optimally redistributes across investors. But then there are capital gains driven entirely by discount rate changes. Or the starting point may be a BGP with a constant discount rate \bar{R} but dividend growth $D_t = G^t D_0$ and then discount rates change to generate capital gains. As expected, both in Proposition 5 and the particular implementation (30), all terms involving ΔD_t drop and optimal redistributive taxes condition only on realized trades $\{x_t(\theta), X_t\}$ and not on asset holdings $\{k_t(\theta), K_t\}$.

Special Case 2. It is also instructive to consider the second extreme case from Section 1 where asset prices change exclusively because of a change in future dividends. First, note from (8) that the the asset price in period 0 satisfies $p_0 = \sum_{s=1}^T R_{0 \rightarrow s}^{-1} D_s$. Now consider a change in prices and dividends such that

$$\frac{\Delta D_{t+1} + \Delta p_{t+1}}{\Delta p_t} = \frac{\bar{D}_{t+1} + \bar{p}_{t+1}}{\bar{p}_t} \quad \text{for all } t \quad (31)$$

This means that asset prices change in such a way that the rates of return \bar{R}_{t+1} remain unchanged for all $t \geq 1$, so the asset price change is entirely driven by dividend changes:

$$\Delta p_t = \sum_{s=1}^{T-t} \bar{R}_{t \rightarrow t+s}^{-1} \Delta D_{t+s}. \quad (32)$$

Analogous to the example in Figure 7, the economy could initially be in a steady state with constant \bar{D} , \bar{p} and \bar{R} but then there are capital gains driven by a change in future dividends $\{\Delta D_t\}$ but with the discount rate being unaffected.

Then we obtain the following result:

Corollary 3. *Suppose the change in prices $\{\Delta p_t\}_{t=0}^T$ and dividends $\{\Delta D_t\}_{t=0}^T$ satisfies (31). Then optimal taxes are such that*

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \bar{T}_t(\theta) + [k_0(\theta) - \Omega(\theta)K_0] (\Delta D_0 + \Delta p_0)$$

In words, the change in the present value of taxes in this special case is given by the accrued gains in period 0, precisely like in Corollary 2 in the two-period model. One particular implementation is a one-time wealth tax or accrual-based capital gains tax at $t = 0$:

$$T_0(\theta) = \bar{T}_0(\theta) + [k_0(\theta) - \Omega(\theta)K_0] (\Delta D_0 + \Delta p_0).$$

Another useful perspective on Corollary 3 which explains why only initial asset holdings $k_0(\theta)$ show up is via the alternative tax implementation that taxes changes in total capital income due to changes in returns $\{\Delta R_t\}$ already mentioned in Section 3.4. To this end, Proposition 9

in Appendix D.3 spells out the multi-period analogue of Proposition 2. Analogous to (30), a natural per-period implementation is

$$T_t(\theta) = \bar{T}_t(\theta) + a_t(\theta)\Delta R_t - \Omega(\theta)A_t\Delta R_t \quad \text{for all } t.$$

The key observation for understanding why only initial asset holdings $k_0(\theta)$ show up is that Special Case 2 with purely dividend-driven asset prices means that, while the initial return changes $\Delta R_0 = (\Delta D_0 + \Delta p_0)/p_{-1} \neq 0$, all returns going forward are unchanged $\Delta R_t = 0$ for all $t \geq 1$. Hence only the time-zero terms $a_0(\theta)\Delta R_0$ show up in the tax formula. But, since $a_0(\theta) = p_{-1}k_0(\theta)$, this exactly equals the term $k_0(\theta)(\Delta D_0 + \Delta p_0)$ in Corollary 3. In contrast, in all other cases except Special Case 2, rising asset prices affect returns going forward so that $\Delta R_t \neq 0$ for $t \geq 1$ and hence there are extra terms in the optimal tax formula – see Section 3.4 for the intuition.

However, even in Special Case 2 with a purely dividend-driven asset price change, this Haig-Simons tax only works once in the initial period, not each period. Indeed, from (32) we have $\Delta p_0 = \sum_{t=1}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta D_t$ and hence another way of writing the tax change is

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \bar{T}_t(\theta) + [k_0(\theta) - \Omega(\theta)K_0] \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta D_t.$$

Hence, a period-by-period implementation would set

$$T_t(\theta) = \bar{T}_t(\theta) + k_0(\theta)\Delta D_t - \Omega(\theta)K_0\Delta D_t \quad \text{for all } t,$$

which does not correspond to a tax on the accrued gains (nor dividend income) in each period.

Deferral advantage. As already noted, the literature on capital gains taxation has emphasized that a realization-based tax system can create a lock-in effect due to the interest advantage from deferring to liquidate appreciated assets. While an accrual-based system as under the Haig-Simons comprehensive tax base would avoid this issue, Vickrey (1939) and Auerbach (1991) have proposed modified (“retrospective”) realization-based capital gains taxes that eliminate this deferral advantage. These tax systems make the tax rate contingent on the holding period, effectively charging interest on delayed realizations.²⁷

The optimal tax policy described in Proposition 5 shares precisely this feature by fixing the present-value of taxes. In particular, one implementation of the optimum would be to set

$$\Delta T_0 = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta)\Delta p_t + k_t(\theta)\Delta D_t - \Omega(\theta)(X_t\Delta p_t + K_t\Delta D_t)]$$

²⁷For instance, Vickrey (1939) writes of a desirable tax system that “The discounted value of the series of tax payments made by any taxpayer should be independent of the way in which his income is allocated to the various income years.”

and $\Delta T_t = 0$ all $t \neq 0$. An alternative would be to set

$$\Delta T_1 / \bar{R}_1 = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Omega(\theta)(X_t \Delta p_t + K_t \Delta D_t)]$$

and $\Delta T_t = 0$ all $t \neq 1$. Therefore,

$$\Delta T_1 = \bar{R}_1 \times \Delta T_0,$$

i.e. the tax liability is exactly growing at the rate of interest. While this connection to the Vickrey and Auerbach proposals is interesting, a full analysis of the lock-in effect would require studying optimal distortive taxation with portfolio choice (see Section 4.4).

Baseline prices and dividends. As the multi-period model makes particularly clear, the tax on the trading gains and losses $x_t(\theta) \Delta p_t$ is computed based on some price change each period that compares the (new) price p_t to some (old) reference price \bar{p}_t . As in Figure 7, this reference price could be the price in some initial steady state $\bar{p}_t = \bar{p}$. Alternatively, it could be the price path on an initial BGP on which dividends and hence prices grow at a constant rate $\bar{p}_t = G^t \bar{p}_0$.

As a result, the reference price does not necessarily coincide with the historical basis at which the investor purchased the asset in some previous period, so the trading gains and losses do not generally correspond to realized capital gains in practice. Instead, the tax is designed to compensate individuals for asset price changes relative to some benchmark that needs to be defined, including for investors who purchase the asset and have not owned it in the past. For example, consider again the BGP example where dividends and prices initially grow at a constant rate. Then suppose there is a change (such as a fall in interest rates) that leads to a different path of prices and dividends. If the redistributive goal is to compensate the winners and losers from this change, the tax system is adjusted based on a comparison, each period, between the actual prices and dividends and the ones that would have occurred if the balanced growth path had continued.

6 Extensions

In this section, we show how the results derived in the benchmark setting extend to richer environments. In particular, we discuss extensions with a closed economy general equilibrium model, heterogeneous returns, aggregate risk and borrowing, and intergenerational transfers. For simplicity, we return to the two-period case.

6.1 General equilibrium

Our baseline model features a small open economy with an exogenously given asset price and dividend. Instead, we now consider a closed economy with the asset in fixed supply, so

$$\int k_0(\theta) dF(\theta) = \int k_1(\theta) dF(\theta) = K. \quad (33)$$

The asset price p must adjust to satisfy the market clearing condition (33). If preferences are given by (6), the equilibrium asset price p^* can be solved in closed form:

$$p^* = \beta D \left(\frac{Y_0}{Y_1 + DK} \right)^{\frac{1}{\sigma}}. \quad (34)$$

This illustrates the various potential drivers of asset price changes in general equilibrium. A particularly natural one is an increase in the discount factor β , which, holding dividends fixed, increases the asset price p^* proportionally. This allows us to obtain the following result:

Proposition 6. *Suppose the equilibrium asset price changes from \bar{p}^* to $p^* = \bar{p}^* + \Delta p^*$, holding dividends and the aggregate endowment fixed. Then the optimal tax $T_0(\theta)$ is given by*

$$T_0(\theta) = \bar{T}_0(\theta) + \bar{x}(\theta)\Delta p^* = \bar{T}_0(\theta) + x(\theta)\Delta p^*$$

where $\bar{T}_0(\theta)$ is the optimal tax at the initial price \bar{p}^* and $\bar{x}(\theta) = x(\theta)$ are investor θ 's asset sales at the initial and old prices \bar{p}^* and p^* , respectively.

Hence, we obtain the same result as in Corollary 1 except that, since aggregate asset sales X must be zero in the closed economy, the intercept term vanishes. Moreover, individual asset sales in fact remain unchanged in response to the asset price change, so $x(\theta) = \bar{x}(\theta)$ for all θ . Intuitively, in the closed economy, total resources do not change when we hold dividends and the aggregate endowment fixed. Hence, the Pareto planner aims to get each investor back to its original consumption bundle after the asset price change. The tax reform in Proposition 6 achieves this through the Slutsky compensation as in Lemma 1.

We can also consider a change in dividends D in general equilibrium. Equation (34) reveals that an increase in D has a less than proportional effect on the asset price p^* due to the indirect effect on the aggregate endowment. Hence, a change in dividends will simultaneously increase the equilibrium rate of return $R^* = D/p^*$. As a result, the knife-edge result in Corollary 2, which lended support to a Haig-Simons accrual-based tax in the special case of a purely dividend-driven asset price change, does not extend to general equilibrium.

6.2 Heterogeneous returns

So far, we have assumed that investors are heterogeneous in their initial endowments $k_0(\theta)$ and incomes $y_0(\theta)$ and $y_1(\theta)$, but they all achieve the same dividends D per unit of their asset holdings $k_1(\theta)$ in period 1. We next show how our results extend to the case with heterogeneous dividends $D(\theta)$ which implies that different investors earn different returns $R(\theta) \equiv D(\theta)/p$.²⁸

Just introducing this additional heterogeneity into our baseline model does not change our results on Pareto optimal tax policy. The reason is that it is efficient, in the absence of further frictions, to allocate all asset holdings to the individual with the highest dividends $D^{\max} \equiv \max_{\theta} D(\theta)$. Hence, the planner can transfer resources at rate of return $R = D^{\max}/p$,

²⁸See, e.g., Gerritsen et al. (2020), Schulz (2021) and Guvenen et al. (2023, 2024) for recent analyses of capital taxation with heterogeneous returns.

returning us effectively to the case without return heterogeneity. Those individuals with lower dividends will not hold the asset, but the government saves for them (using the highest-return individual) through taxes and transfers $T_0(\theta)$ and $T_1(\theta)$. Hence, Proposition 1 goes through, with the only twist that almost all investors will be sellers with $x(\theta) = k_0(\theta)$ and $k_1(\theta) = 0$.

Trading with adjustment costs. To prevent this trivial outcome, it is useful to introduce a second asset as well as some trading friction. Individuals can save both in a bond $b(\theta)$, which trades at price q in period 0, and in the asset, which trades at price p in period 0 and delivers dividends $D(\theta)$ in period 1. Thus, an investor's budget constraints are:

$$\begin{aligned} c_0(\theta) + qb(\theta) &= px(\theta) - \chi(x(\theta)) + y_0(\theta) - T_0(\theta) \\ c_1(\theta) &= D(\theta)(k_0(\theta) - x(\theta)) + b(\theta) + y_1(\theta). \end{aligned}$$

The trading frictions are captured by an adjustment cost $\chi(x)$ that is increasing in $|x|$ and convex. This cost ensures that it is no longer efficient to allocate all capital to the individual with the highest return.

The present-value budget constraint is:

$$c_0(\theta) + qc_1(\theta) = qD(\theta)(k_0(\theta) - x(\theta)) + px(\theta) - \chi(x(\theta)) + y_0(\theta) + qy_1(\theta) - T_0(\theta),$$

which immediately implies that an investor's optimal asset sales $x(\theta)$ satisfy

$$qD(\theta) + \chi'(x(\theta)) = p \tag{35}$$

The left-hand side captures the marginal cost of selling more assets: the investor will have less dividend income and will need to pay the additional trading cost. On the other hand, the asset price on the right-hand side is the additional revenue from the sale. As a result, due to the convex adjustment cost, investors with higher returns $D(\theta)$ will sell less and therefore hold more of the asset. Also note that the presence of adjustment costs in (35) implies that

$$R(\theta) \equiv \frac{D(\theta)}{p} \geq \frac{1}{q}, \tag{36}$$

so that (i) the usual no-arbitrage condition equalizing the return on the asset $R(\theta)$ to that on the bond $1/q$ may not hold and, therefore, (ii) different investors θ may obtain different asset returns $R(\theta)$ in equilibrium. This opens up the door to the returns $R(\theta)$ of different investors changing differentially in response to heterogeneous cash flow changes, which will be important below.

The aggregate resource constraint can be written as

$$\int c_0(\theta)dF(\theta) + q \int c_1(\theta)dF(\theta) = Y \tag{37}$$

where

$$Y = Y_0 + qY_1 + \max_{\{x(\theta)\}} \int [px(\theta) + qD(\theta)(k_0(\theta) - x(\theta)) - \chi(x(\theta))] dF(\theta)$$

Thus, the first-best Pareto problem takes the same form as in Section 3. For simplicity, we continue to assume our preference specification (6) and normalize $T_1(\theta) = 0$, which leads to the following result:

Proposition 7. *Suppose the equilibrium asset price changes from \bar{p} to $p = \bar{p} + \Delta p$, holding dividends $D(\theta)$ and the bond price q fixed. Then the optimal tax $T_0(\theta)$ satisfies the second-order approximation*

$$T_0(\theta) \approx \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p - \frac{1}{2} \left[\chi''(\bar{x}(\theta))\Delta x(\theta)^2 - \Omega(\theta) \int \chi''(\bar{x}(\theta))\Delta x(\theta)^2 dF(\theta) \right]$$

where $\Delta x(\theta) \equiv x(\theta) - \bar{x}(\theta)$.

Hence, even with heterogeneous returns and trading frictions, Corollary 1 goes through to first order, and an additional second-order term emerges that captures the change in adjustment costs due to the asset price change.

General equilibrium. This result is particularly useful when combining it with our previous general-equilibrium analysis. Suppose the asset is in fixed supply, as in the preceding subsection, and assume, for simplicity, a quadratic adjustment cost $\chi(x) = \kappa x^2$. Then the optimality condition (35) together with the market clearing condition $X = 0$ immediately implies

$$p^* = q \int D(\theta) dF(\theta).$$

The equilibrium price equals the discounted average dividends in the economy. This is intuitive, since even investors with a low dividend can sell their asset to other investors with higher dividends, so the asset price must reflect the average dividend when χ is symmetric around zero and quadratic. As already anticipated above, in equilibrium, different investors experience differential returns given by

$$R^*(\theta) \equiv \frac{D(\theta)}{p^*} = \frac{D(\theta)}{q \int D(\bar{\theta}) dF(\bar{\theta})}.$$

Consider now an increase in dividends $D(\theta)$ for *some subset* of the investors in the economy. This will induce an increase in the equilibrium asset price p^* for *all* investors, including for those whose dividends did not change.²⁹ Put differently, those investors whose dividends $D(\theta)$ increase experience an *increase* in their asset return $R^*(\theta)$; however those investors whose dividends $D(\theta)$ remain unchanged, actually experience a *decline* in their asset return $R^*(\theta)$. Since these latter investors face a pure asset price increase without a simultaneous dividend change, their optimal tax change, to a first order, is given by $\Delta T_0(\theta) \approx x(\theta)\Delta p^*$ as in Propositions 6 and 7.

²⁹If bonds are in fixed supply as well, then the bond price q will also adjust, but in general this will not undo the change in the average dividends.

Importantly, this is true even though the asset price change is ultimately driven by a dividend change. Hence, our knife-edge result from Corollary 2, which found that a Haig-Simons accrual-based tax can implement the optimum in the special case of a purely dividend-driven asset price change, does not survive in this richer model. The reason is that Corollary 2 relied on the fact that the effect of the asset price change and the dividend change happened to cancel in the baseline model where everyone achieves the same rate of return. With heterogeneous dividends, these effects will no longer cancel in general (not even for the investors whose dividends change) so the Haig-Simons tax never applies. By contrast, a tax on both realized capital gains and dividends as in Proposition 1 continues to work.

6.3 Aggregate Risk

Consider our benchmark environment, but extended to allow for aggregate risk in the second-period dividend. In particular, let $D(s)$ denote the realized dividend in the second period, where $s \in S$ denotes the aggregate state of the economy in period $t = 1$ and S denotes the finite set of possible outcomes. Let $\pi(s)$ denote the probability state s will be realized.

Preferences. We extend consumer preferences to incorporate risk aversion following Epstein and Zin (1989). In particular, let μ be the following certainty equivalence aggregator:

$$\mu(\{c_1(s)\}_{s \in S}) \equiv \left(\sum_{s \in S} \pi(s) c_1(s)^{1-\alpha} \right)^{\frac{1}{1-\alpha}},$$

where $\alpha \geq 0$ is the coefficient of relative risk aversion. With this, we can extend the preferences defined in (6):

$$U(c_0, \{c_1(s)\}) = G(C(c_0, \mu(\{c_1(s)\}))),$$

where, as in (6), $G = C^{1-\gamma}/(1-\gamma)$ and C is a CES aggregator over c_0 and μ with elasticity σ .

Asset Markets. Aggregate dividend risk implies that saving in capital is risky. Just like in the preceding subsection, we therefore allow the agents to also trade in a risk-free bond. Let p denote the period-0 price of capital and q denote the price of the zero-coupon risk-free bond. The consumer's flow budget constraints become:

$$\begin{aligned} c_0(\theta) &= p(k_0(\theta) - k_1(\theta)) - qb(\theta) + y_0(\theta) - T_0(\theta) \\ c_1(\theta, s) &= D(s)k_1(\theta) + b(\theta) + y_1(\theta) - T_1(\theta, s), \quad \forall s \in S, \end{aligned}$$

where b denotes the quantity of the risk-free bond purchased (which can be negative when the agent borrows). Note that we allow taxes and transfers in the second period to be indexed by s .

We retain the small-open-economy assumption, and denote $M(s)$ as the stochastic discount factor of the representative counterparty in global financial markets. In particular, the period-0 Arrow-Debreu price of a unit of consumption delivered in state s is $\pi(s)M(s)$. No arbitrage

implies:

$$\begin{aligned} p &= \sum_{s \in S} \pi(s) M(s) D(s) \\ q &= \sum_{s \in S} \pi(s) M(s). \end{aligned} \tag{38}$$

In order to ensure that aggregate risk is relevant, we assume

$$\int T_0(\theta) dF(\theta) = \int T_1(\theta, s) dF(\theta) = 0 \quad \forall s \in S,$$

so the economy cannot insure itself other than through trading capital and the risk-free bond with the rest of the world. In other words, it does not have access to the full set of Arrow-Debreu insurance markets. Hence, aggregate second-period allocations are spanned by the risk-free bond and payoffs to capital:

$$C_1(s) \equiv \int c_1(\theta, s) dF(\theta) = D(s)K_1 + Y_1 + B, \quad \forall s \in S, \tag{39}$$

where $K_1 = \int k_1(\theta) dF(\theta)$, $Y_1 = \int y_1(\theta) dF(\theta)$, and $B = \int b(\theta) dF(\theta)$.³⁰

First Best. The first-best allocation is the solution to

$$\max_{\{c_0(\theta), \{c_1(\theta, s)\}, X\}} \mathbb{E} \int \omega(\theta) U(c_0(\theta), \mu(\{c_1(\theta, s)\})) dF(\theta) \tag{40}$$

subject to

$$\int c_0(\theta) dF(\theta) + q \int c_1(\theta, s) dF(\theta) = Y_0 + qY_1 + pX + qD(s)(K_0 - X) \quad \forall s \in S.$$

As shown in Appendix E.4, the homotheticity of preferences implies that the first-best allocation can be studied as a two-stage budgeting problem. The first stage is to solve a single-agent portfolio problem involving aggregates C_0 and $C_1(s)$:

$$\max_{\{C_0, \{C_1(s)\}, X\}} U(C_0, \mu(\{C_1(s)\})) \quad \text{s.t.} \quad C_0 + qC_1(s) = Y_0 + qY_1 + pX + qD(s)(K_0 - X).$$

Letting an asterik denote the solution to the first stage, the second stage involves allocating aggregate consumption according to the rule:

$$c_0^*(\theta) = \Omega(\theta)C_0^*, \quad c_1^*(\theta, s) = \Omega(\theta)C_1^*(s), \tag{41}$$

where $\Omega(\theta)$ is defined in (21).

Decentralization of First Best. As in the benchmark environment, the fact that individuals can trade assets generates an indeterminacy in the tax system that decentralizes the first-best

³⁰Without loss of generality, we assume that the government does not own bonds or capital directly, and thus the aggregates B and K_1 reflect the aggregated holdings of private individuals.

allocation. With two assets, there are two dimensions of indeterminacy, spanned by the payoffs to the risk-free bond and risky capital. Specifically:

Lemma 3. *Suppose $a : \Theta \rightarrow \mathbb{R}$ and $m : \Theta \rightarrow \mathbb{R}$ are functions of θ such that $\int a(\theta)dF = \int m(\theta)dF = 0$. Then there exists a first-period tax schedule $T_0 : \Theta \rightarrow \mathbb{R}$ that implements the first-best when combined with the following second-period tax schedules:*

$$T_1(\theta, s) = a(\theta) + m(\theta)D(s) \quad \forall \theta, s.$$

That is, the tax schedule can be an arbitrary linear function of the payoffs to risky capital. This follows from the fact that individuals can always adjust their private holding of bonds and capital to account for differences in the tax system that are spanned by the payoffs to bonds and capital.

Comparative Statics. We are now in position to revisit how changes in asset prices and cash flows induce changes in the optimal tax schedules. The natural extension of the “discount rate shock” in Corollary 1 to this environment is a change in the pricing kernel M , holding D fixed. That is, $M(s) = \bar{M}(s) + \Delta M(s)$. For example, suppose attitudes toward risk change, or the time discounting inherent in M changes. This would induce changes in the price of capital and risk-free bonds even in the absence of changes to dividends. Additionally, it could be the case that cash flows themselves change: $D(s) = \bar{D}(s) + \Delta D(s)$. The next result considers the general case and extends Proposition 1 to the case with risky capital.

Proposition 8. *Suppose the pricing kernel changes from $\bar{M}(s)$ to $M(s) = \bar{M}(s) + \Delta M(s)$ and dividends from $\bar{D}(s)$ to $D(s) = \bar{D}(s) + \Delta D(s)$. Let $\Delta p \equiv \sum_s \pi(s)M(s)D(s) - \sum_s \pi(s)\bar{M}(s)\bar{D}(s)$ and $\Delta q \equiv \sum_s \pi(s)\Delta M(s)$. Then the following change to the tax schedules is an optimal response:*

$$\begin{aligned} T_0(\theta) &= \bar{T}_0(\theta) + x(\theta)\Delta p - b(\theta)\Delta q - \Omega(\theta) [X\Delta p - B\Delta q] \\ T_1(\theta, s) &= \bar{T}_1(\theta, s) + k_1(\theta)\Delta D(s) - \Omega(\theta)K_1\Delta D(s). \end{aligned}$$

As before, it is easiest to discuss this result separating changes in discount factors from cash flows. For the case of an asset-price change exclusively due to discount-rate variation (Special Case 1), the counterpart to Corollary 1 is:

Corollary 4. *Suppose the pricing kernel changes $\Delta M(s) \neq 0$ for some $s \in S$, while dividends are unchanged: $\Delta D(s) = 0$ for all $s \in S$. Then an optimal change in the tax schedule is to maintain second period taxes $T_1(\theta, s) = \bar{T}_1(\theta, s)$, $\forall s, \theta$, and set*

$$T_0(\theta) = \bar{T}_0(\theta) + x(\theta)\Delta p - b(\theta)\Delta q - \Omega(\theta) [X\Delta p - B\Delta q].$$

Compared to Corollary 1, the only difference is the additional compensation for changes in the bond price Δq . The intuition is simply that a change in the interest rate on the bond redistributes between borrowers and savers, and the first-best tax response counteracts this.

A noteworthy example is a change in the pricing kernel $\Delta M(s)$ such that $\sum_{s \in S} \pi(s) \Delta M(s) = 0$. Then $\Delta q = 0$, so the risk-free rate is unchanged, and since we also hold dividends fixed, this corresponds to a pure risk-premium change. In this case, we collapse back to Corollary 1. In other words, even in this richer setting, the optimal tax response to an asset price change induced by a risk-premium change targets the realized trading gains and losses exactly like in our deterministic benchmark model.

For the case of asset price changes driven by only cash-flow changes (Special Case 2), we show that Corollary 2 from the deterministic model also goes through:

Corollary 5. *Suppose $\Delta D(s) \neq 0$ for some $s \in S$, while $\Delta M(s) = 0$ for all $s \in S$. Then an optimal change in the tax schedule is to keep $T_1(\theta, s) = \bar{T}_1(\theta, s)$ unchanged and set*

$$T_0(\theta) = \bar{T}_0(\theta) + k_0(\theta) \Delta p - \Omega(\theta) K_0 \Delta p.$$

Hence, a Haig-Simons tax on unrealized gains can be applied in period 0 in this case, but, as before, it is knife-edge and breaks down whenever the pricing kernel changes.

Subjective beliefs. Subjective beliefs—optimism and pessimism—may play a role in driving asset prices. For example, Adam et al. (2017) and Bordalo et al. (2023) argue that, while measures of actual current and expected future cash flows cannot account for asset-price fluctuations, variation in subjective expectations about these cash flows can. The setup in this section allows us to consider this case. To this end, recall from (38) that the asset prices p and q depend on the probabilities $\pi(s)$ with which different states $s \in S$ occur as well as the stochastic discount factor $M(s)$ with which a representative counterparty in global financial markets discounts these states. To capture subjective beliefs, we can replace the actual probabilities $\pi(s)$ by some alternative subjective beliefs $\tilde{\pi}(s)$, while assuming that the social planner continues to use the correct probabilities $\pi(s)$. More optimism in financial markets then corresponds to $\tilde{\pi}(s)$ changing in such a way as to put a higher probability on states s in which cash flows $D(s)$ are high. Modeled in this way, optimism about cash flows generates an increase in the asset price p while leaving actual cash flows unaffected. Changes in subjective beliefs are therefore equivalent to discount rate shocks.

6.4 Borrowing versus Selling

An argument that frequently comes up in discussions about the redistributive effects of asset-price changes is that wealthy individuals do not necessarily need to sell their appreciated assets; instead they can just borrow against the appreciated asset values. The Economist (2024) provides an instructive example:

“Say you own a successful business – so successful that your stake in it is worth \$1bn. How should you finance your spending? If you [...] sell \$20m-worth of shares [...], the entire sum represents capital gains and will be taxed at 20%, which would mean a \$4m hit. What if, instead, you called up your wealth manager and agreed to put up \$100m-worth of equity as collateral for a \$20m loan. In 2021

the interest rate on the loan might have been just 2% a year, meaning that returns from holding the equity, rather than selling it, would easily have covered the cost of servicing the borrowing. Because the proceeds of loans, which must be eventually repaid, are not considered income, doing so would have incurred no tax liability at all.”

Both the setups in Sections 6.2 (heterogeneous returns) and 6.3 (aggregate risk) feature borrowing and are therefore useful for analyzing how optimal taxation should treat borrowing versus selling. The most instructive case is that of positive capital gains $\Delta p > 0$ but without a corresponding change in interest rates $\Delta q = 0$. In Section 6.3 with risk, this case obtains for a pure risk-premium change – see equation (38). Setting $\Delta q = 0$ in Corollary 4 yields

$$T_0(\theta) = \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p.$$

Similarly, in Section 6.2 with heterogeneous returns, Proposition 7 shows that the same expression holds to first order. Perhaps surprisingly, this result shows that the tax formula is *independent of* whether and how much investors borrow when their assets appreciate and is, in fact, identical to that in Corollary 1. An analogous result can be shown in a version of the multi-period model in Section 5 with borrowing.

The intuition is that, also with the option to borrow, investors need to sell their appreciating assets *at some point* in order to benefit from rising asset prices (the exception is “stepped-up basis” which we discuss in the next subsection). The key observation (which is sometimes missed in the popular debate) is that loans need to be repaid at some point. If investors never sell their assets, they will need to repay their loans out of income they could have otherwise consumed and hence they do not benefit from the capital gains.³¹ On the other hand, if investors do sell to repay the loan, the realized trade should be taxed at that point.

The example at the beginning of this subsection emphasized an important motive for borrowing rather than selling an asset: the asset’s return often exceeds the rate at which investors can borrow. While the setup with aggregate risk in Section 6.3 does not allow for this possibility, the setup with heterogeneous returns in Section 6.2 has precisely this feature – see equation (36). Still, this consideration does not change the optimal tax formula. The intuition is that, while such return differences are undoubtedly important, they are not special to the case of wealthy individuals borrowing against appreciating assets. Instead they are a feature of *any* levered investment strategy. For example, many homeowners with an outstanding mortgage invest some of their income in the stock market rather than pre-paying their mortgage, precisely because stock returns exceed mortgage interest rates. Investors using levered investment strategies to take advantage of such return differences is simply an orthogonal issue that should not be considered tax avoidance.

³¹A case in which things are more complicated and investors may benefit from capital gains even without selling is the case of a binding collateral constraint that gets relaxed by a rising asset price. [Fagereng et al. \(2023\)](#) discuss this case at length. Given our focus on the top of the wealth distribution, we here abstract from binding constraints.

6.5 Bequests and Suboptimality of Step-Up in Basis at Death

We finally consider a version of our model with multiple generations in which parents bequeath to their children. We use this version to consider a peculiarity of the tax system in the U.S. and many other advanced economies: step-up in basis at death for inherited assets, a tax rule that eliminates the taxable capital gain that occurred between the original purchase of the asset and the heir's acquisition, thereby reducing the heir's tax liability.³²

To keep things simple, we model dynasties of non-overlapping generations that are altruistic toward their offspring à la [Barro and Becker \(1989\)](#). A new generation of investors is born every τ years and investors live for $\tau - 1$ periods. An investor of dynasty θ born at time t (cohort t) has lifetime utility

$$V_t(\theta) = U(c_t(\theta), \dots, c_{t+\tau-1}(\theta)) + \alpha\beta^\tau V_{t+\tau}(\theta), \quad (42)$$

where $0 \leq \alpha \leq 1$ measures altruism toward the next generation and $U(c_t, \dots, c_{t+\tau-1})$ is given by (6) with $T = \tau - 1$. The sequential budget constraint is still given by (2) but now with the convention that $k_\tau(\theta), k_{2\tau}(\theta), k_{3\tau}(\theta)$, and so on denote bequests left by investors in their last year of life toward their offspring in the next generation. Because investors are altruistic, these bequests will generally be positive. As is standard, the Barro-Becker assumption implies that we can work with the preferences of dynasties.³³ The Pareto problem is to maximize $\int \omega(\theta) V_0(\theta) dF(\theta)$ subject to (7) where $\omega(\theta)$ is the Pareto weight on dynasty θ .

A particularly simple case is that of full altruism $\alpha = 1$ and $\gamma = 1/\sigma$. Then the time-0 utility of dynasty θ is given by $V_0(\theta) = \sum_{t=0}^{\infty} \beta^t c_t(\theta)^{\frac{\sigma-1}{\sigma}}$, which is a special case of (6) with $T = \infty$. Therefore, everything collapses to the multi-period model from Section 5, in particular Proposition 5 applies. More generally, when $\alpha < 1$ or $\gamma \neq 1/\sigma$, Proposition 5 is modified to include a time-varying consumption allocation rule $\Omega_t(\theta)$ in place of the constant rule $\Omega(\theta)$ (i.e. the planner now allocates consumption $c_t(\theta) = \Omega_t(\theta)C_t$ to dynasty θ).

Suboptimality of Step-Up of Basis on Death. Because (a modified) Proposition 5 still applies, so does the discussion about the baseline relative to which capital gains are calculated. As discussed there, a natural benchmark is the price path on an initial BGP on which dividends and hence prices grow at a constant rate $\bar{p}_t = G^t \bar{p}_0$. Step-up of basis on death would instead correspond to a case in which the baseline price \bar{p}_t resets to the current market price p_t every τ years, i.e. whenever a generation dies and is replaced by the next generation. In the Barro-Becker setup analyzed here, from the point of a view of a dynasty or the social planner,

³²Many good explanations of step-up in basis can be found on the internet, particularly by financial and estate planning services. Some of these are explicit that they consider the rule to be a loophole, for example [Trust and Will \(2024\)](#) which begins the discussion thus: "Loopholes – you may not always use them, but when you do need them, you're sure glad they're there. [...] The Step-Up in Basis loophole is used to circumvent capital gains taxes, or to pay the least amount of this type of inheritance tax as is legally possible."

³³For example, repeated substitution of (42) implies that the dynasty θ 's utility at time 0 is given by

$$V_0(\theta) = U(c_0(\theta), \dots, c_{\tau-1}(\theta)) + \alpha\beta^\tau U(c_\tau(\theta), \dots, c_{2\tau-1}(\theta)) + \alpha^2\beta^{2\tau} U(c_{2\tau}(\theta), \dots, c_{3\tau-1}(\theta)) + \dots$$

there is nothing special about the dates at which one generation passes the baton to the next and therefore also no argument for resetting the basis in this way. Instead, a natural approach is the “carry-over basis” already used by a number of countries including Germany, Italy, and Japan (OECD, 2021).

Buy, Borrow, Die. In recent years, a tax avoidance strategy known as “buy, borrow, die” has become popular among wealthy families, particularly in the U.S. (e.g. Ensign and Rubin, 2021; The Economist, 2024; Fox and Liscow, 2024). The idea is to borrow against appreciating assets rather than selling them and then taking advantage of the stepped-up basis at death, thereby allowing those who employ this strategy to avoid paying capital-gains taxes altogether. In combination with the discussion in the preceding subsection on borrowing versus selling, our results suggests that the stepped-up basis loophole should be eliminated. Absent stepped-up basis, the wealthy would still benefit from borrowing against high-return assets with lower-interest loans as described in the preceding subsection. But this is just like any other levered investment and should not be considered a tax avoidance strategy.

7 Conclusion

We “put the ‘finance’ into ‘public finance’,” meaning that we study optimal redistributive taxation with changing asset prices. Importantly, we adopt the modern finance view that asset prices change not only because of changing cash flows but also due to changes in discount rates, risk premia, or subjective beliefs.

To summarize our results, it is useful to juxtapose them with the following naïve intuition implicit in proposals for wealth taxes or taxes on unrealized capital gains: when the value of Jeff Bezos’ Amazon stocks doubles so should his tax liability. Our results show that this intuition is, in general, incorrect. Optimal taxes instead generally depend on (i) whether Bezos sells his Amazon shares and (ii) whether and by how much cash flows, here Amazon’s profits, increase. In our baseline model these are, in fact, the only determinants of optimal taxes. Generalizations of the type considered in Section 6 complicate the optimal tax formulas in some cases but it remains true that taxes that target only asset holdings rather than asset transactions are generally suboptimal.

References

- Adam, Klaus, Albert Marcet, and Johannes Beutel, “Stock Price Booms and Expected Capital Gains,” *American Economic Review*, 2017, 107 (8), 2352–2408.
- Agersnap, Ole and Owen Zidar, “The Tax Elasticity of Capital Gains and Revenue-Maximizing Rates,” *American Economic Review: Insights*, 2021, 3 (4), 399–416.
- Atkeson, Andrew, Jonathan Heathcote, and Fabrizio Perri, “There is No Excess Volatility Puzzle,” *NBER Working Paper 32481*, 2024.
- Atkinson, Anothony and Joseph Stiglitz, “The design of tax structure: Direct versus indirect taxation,” *Journal of Public Economics*, 1976, 6, 55–75.

- Auerbach, Alan**, “Capital Gains Taxation and Tax Reform,” *National Tax Journal*, 1989, 42 (3), 391–401.
- , “Retrospective Capital Gains Taxation,” *American Economic Review*, 1991, 81 (1), 167–178.
- Banks, James and Peter Diamond**, “The base for direct taxation,” in Stuart Adam, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba, eds., *Dimensions of Tax Design: The Mirrlees Review*, Oxford University Press, 2010, pp. 548–648.
- Barro, Robert and Gary Becker**, “Fertility Choice in a Model of Economic Growth,” *Econometrica*, 1989, 57 (2), 481–501.
- Bastani, Spencer and Daniel Waldenstrom**, “How Should Capital Be Taxed?,” *Journal of Economic Surveys*, 2020, 34 (4), 812–846.
- Benhabib, Jess and Bálint Szöke**, “Optimal Positive Capital Taxes at Interior Steady States,” *American Economic Journal: Macroeconomics*, 2021, 13 (1), 114–150.
- Bonnet, Odran, Pierre-Henri Bono, Guillaume Flamerie de La Chapelle, and Etienne Wasmer**, “Does housing capital contribute to inequality? A comment on Thomas Piketty’s Capital in the 21st Century,” *SciencePo Working paper 2014-07*, 2014.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, Matthew OBrien, and Andrei Shleifer**, “Long-Term Expectations and Aggregate Fluctuations,” *NBER Macroeconomics Annual*, 2023, 38.
- Campbell, John**, *Financial Decisions and Markets*, Princeton University Press, 2018.
- and **Robert Shiller**, “The dividend-price ratio and expectations of future dividends and discount factors,” *The Review of Financial Studies*, 1988, 1 (3), 195–228.
- Caramp, Nicolas and Dejanir Silva**, “Monetary Policy and Wealth Effects: The Role of Risk and Heterogeneity,” *UC Davis Working Paper*, 2023.
- Catherine, Sylvain, Max Miller, and Natasha Sarin**, “Social security and trends in wealth inequality,” *Working Paper*, 2020.
- , —, **James Paron, and Natasha Sarin**, “Interest-rate Risk and Household Portfolios,” Technical Report, Working Paper 2024.
- Chamley, Christophe**, “Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives,” *Econometrica*, 1986, 54 (3), 607–622.
- Chari, V. V., Mikhail Golosov, and Aleh Tsyvinski**, “Business Start-ups, The Lock-in Effect, and Capital Gains Taxation,” Technical Report, Working Paper 2005.
- Cioffi, Riccardo**, “Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality,” *Princeton University Working Paper*, 2021.
- Cochrane, John**, “Presidential Address: Discount Rates,” *Journal of Finance*, 2011, 66 (4), 1047–1108.
- Constantinides, George**, “Capital Market Equilibrium with Personal Tax,” *Econometrica*, 1983, 51 (3), 611–636.
- Ensign, Rachel Louise and Richard Rubin**, “Buy, Borrow, Die: How Rich Americans Live Off Their Paper Wealth,” *Wall Street Journal*, 2021.
- Epstein, Larry and Stanley Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 1989, 57 (4), 937.
- Fagereng, Andreas, Matthieu Gomez, Emilien Guin-Bonenfant, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik**, “Asset-Price Redistribution,” *Working Paper*, 2023.
- Farhi, Emmanuel, Christopher Sleet, Iván Werning, and Sevin Yeltekin**, “Non-linear Capital Taxation without Commitment,” *Review of Economic Studies*, 2012, 79, 1469–1493.

- Feldstein, Martin, Joel Slemrod, and Shlomo Yitzhaki**, “The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains,” *The Quarterly Journal of Economics*, 1980, 94 (4), 777–791.
- Ferey, Antoine, Ben Lockwood, and Dmitry Taubinsky**, “Sufficient Statistics for Nonlinear Tax Systems with General Across-Income Heterogeneity,” *American Economic Review*, 2024, *forthcoming*.
- Fox, Edward and Zachary Liscow**, “No More Tax-Free Lunch for Billionaires: Closing the Borrowing Loophole,” *Tax Notes Federal*, 2024.
- Gerritsen, Aart, Bas Jacobs, Alexandra Victoria Rusu, and Kevin Spiritus**, “Optimal Taxation of Capital Income with Heterogeneous Rates of Return,” *CESifo working paper 8395*, 2020.
- Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning**, “New Dynamic Public Finance: A User’s Guide,” *NBER Macroeconomics Annual*, 2007, 21, 317–388.
- Gomez, Matthieu**, “Asset prices and wealth inequality,” *Working paper*, 2016.
- **and Emilien Gouin-Bonenfant**, “A q-theory of inequality,” *Working paper*, 2020.
- Gordon, Myron and Eli Shapiro**, “Capital equipment analysis: the required rate of profit,” *Management science*, 1956, 3 (1), 102–110.
- Greenwald, Daniel, Martin Lettau, and Sydney Ludvigson**, “How the wealth was won: Factors shares as market fundamentals,” *NBER working paper 25769*, 2019.
- , **Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh**, “Financial and total wealth inequality with declining interest rates,” *NBER working paper 28613*, 2021.
- Güvener, Fatih, Gueorgui Kambourov, Burhan Kuruscu, and Sergio Ocampo**, “Book-Value Wealth Taxation, Capital Income Taxation, and Innovation,” *NBER working paper 32585*, 2024.
- , — , — , — , — , **and Daphne Chen**, “Use It or Lose It: Efficiency and Redistributive Effects of Wealth Taxation,” *The Quarterly Journal of Economics*, 2023, 138 (2), 835–894.
- Haig, Robert**, “The Concept of Income – Economic and Legal Aspects,” in Robert M. Haig, ed., *The Federal Income Tax*, New York: Columbia University Press, 1921, pp. 1–28.
- Hellwig, Christian and Nicolas Werquin**, “Using Consumption Data to Derive Optimal Income and Capital Tax Rates,” *Working Paper*, 2024.
- Jermann, Urban**, “Asset pricing in production economies,” *Journal of Monetary Economics*, 1998, 41 (2), 257–275.
- Judd, Kenneth**, “Redistributive taxation in a simple perfect foresight model,” *Journal of Public Economics*, 1985, 28 (1), 59–83.
- Kaldor, Nicholas**, *An Expenditure Tax*, London and New York, relevant excerpt available at <https://benjaminmoll.com/kaldor/>; Routledge, 1955.
- Kuhn, Moritz, Moritz Schularick, and Ulrike Steins**, “Income and wealth inequality in America, 1949–2016,” *Journal of Political Economy*, 2020, 128 (9), 3469–3519.
- Larrain, Borja and Motohiro Yogo**, “Does firm value move too much to be justified by subsequent changes in cash flow,” *Journal of Financial Economics*, 2008, 87 (1), 200–226.
- Leiserson, Greg and Danny Yagan**, “What Is the Average Federal Individual Income Tax Rate on the Wealthiest Americans?,” *U.S. Council of Economic Advisors, Washington, DC, White House*, 2021.
- Lucas, Robert**, “Supply-Side Economics: An Analytical Review,” *Oxford Economic Papers*, 1990, 42 (2), 293–316.
- Martínez-Toledano, Clara**, “House Price Cycles, Wealth Inequality and Portfolio Reshuffling,” *Working Paper, Imperial College London*, 2022.

- Mas-Colell, Andreu, Michael Whinston, and Jerry Green**, *Microeconomic Theory*, New York: Oxford University Press, 1995.
- Mirrlees, James**, “An Exploration in the Theory of Optimum Income Taxation,” *The Review of Economic Studies*, 1971, 38 (2), 175–208.
- Moll, Benjamin**, “Comment on “Sources of U.S. Wealth Inequality: Past, Present, and Future”,” *NBER Macroeconomics Annual*, 2020, 35.
- OECD**, “Inheritance Taxation in OECD Countries,” *OECD Tax Policy Studies*, 2021.
- Paish, Frank**, “Capital Value and Income,” *Economica*, 1940, 7 (28), 416–418.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman**, “Rethinking capital and wealth taxation,” *Oxford Review of Economic Policy*, 2023, 39 (3), 575–591.
- Poterba, James**, “Taxation, risk-taking, and household portfolio behavior,” in A. J. Auerbach and M. Feldstein, eds., *Handbook of Public Economics*, Vol. 3 of *Handbook of Public Economics*, Elsevier, 2002, chapter 17, pp. 1109–1171.
- Rognlie, Matthew**, “Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?,” *Brookings Papers on Economic Activity*, 2015, 46 (1 (Spring)), 1–69.
- Saez, Emmanuel**, “The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes,” *Journal of Public Economics*, 2002, 83, 217–230.
- **and Stefanie Stantcheva**, “A simpler theory of optimal capital taxation,” *Journal of Public Economics*, 2018, 162 (C), 120–142.
- **, Danny Yagan, and Gabriel Zucman**, “Capital Gains Withholding,” *UC Berkeley working paper*, 2021.
- Scheuer, Florian and Alexander Wolitzky**, “Capital Taxation under Political Constraints,” *American Economic Review*, 2016, 106, 2304–2328.
- **and Joel Slemrod**, “Taxing Our Wealth,” *Journal of Economic Perspectives*, Winter 2021, 35 (1), 207–230.
- Schulz, Karl**, “Redistribution of Return Inequality,” *Working paper*, 2021.
- Shiller, Robert**, “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, June 1981, 71 (3), 421–436.
- Shourideh, Ali**, “Taxation of Wealthy Individuals,” *Working Paper*, 2012.
- Simons, Henry**, *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy*, University of Chicago Press Chicago, Ill., 1938.
- Smith, Kate and Helen Miller**, “It’s all about the base: Taxing business owner-managers,” Technical Report, LSE 2023.
- Stantcheva, Stefanie**, “Dynamic Taxation,” *Annual Review of Economics*, 2020, 12, 801–831.
- Straub, Ludwig and Iván Werning**, “Positive Long-Run Capital Taxation: Chamley-Judd Revisited,” *American Economic Review*, 2020, 110 (1), 86–119.
- The Economist**, “Sweet gains, bro,” *June 22 2024 Issue, Finance and Economics Section*, 2024.
- Trust and Will**, “How the Step-Up in Basis Loophole Works,” Technical Report 2024.
- U.S. Department of the Treasury**, *General Explanations of the Administration’s Fiscal Year 2023 Revenue Proposals*, Department of the Treasury, March 2022.
- U.S. Office of Management and Budget**, *Budget of the US Government, Fiscal Year 2023*, US Government Publishing Office, 2022.

- van Binsbergen, Jules**, "Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility," *NBER Working Paper* 27367, 2020.
- Vickrey, William**, "Averaging of Income for Income-Tax Purposes," *Journal of Political Economy*, 1939, 47 (3), 379–379.
- von Schanz, Georg**, "Der Einkommensbegriff und die Einkommensteuergesetze," *FinanzArchiv / Public Finance Analysis*, 1896, 13 (1), 1–87.
- Werning, Iván**, "Optimal Fiscal Policy with Redistribution," *The Quarterly Journal of Economics*, 2007, 122 (3), 925–967.
- Whalley, John**, "Capital Gains Taxation And Interest Rate Changes: An Extension Of Paish's Argument," *National Tax Journal*, 1979, 32 (1), 87–91.
- Wolff, Edward**, "The Stock Market and the Evolution of Top Wealth Shares in the United States," *The Journal of Economic Inequality*, 2022, pp. 1–14.
- Zucman, Gabriel**, "A blueprint for a coordinated minimum effective taxation standard for ultra-high-net-worth individuals," Technical Report, available at <https://gabriel-zucman.eu/files/report-g20.pdf> 2024.

Appendix

A Appendix for Section 1

A.1 Mapping the Neoclassical Growth Model into Our Setup

As is standard, a representative consumer has preferences $\sum_{t=0}^{\infty} \beta^t U(C_t)$ where C_t is consumption, the consumption good is produced according to a constant-returns technology $\mathcal{Y}_t = f(K_t, A_t L_t)$ where K_t is capital, A_t is productivity and L_t is labor, labor is supplied inelastically $L_t = 1$, and the resource constraint is

$$C_t + I_t = \mathcal{Y}_t, \quad K_{t+1} = I_t + (1 - \delta)K_t, \quad (43)$$

where I_t is investment. Importantly, the fact that the consumption good \mathcal{Y}_t can be converted into investment one-for-one immediately pins down the unit price of capital (relative to consumption) at one. We discuss this property in more detail momentarily where we also discuss how to break it.

A.1.1 The Price of Capital in Variants of the Growth Model

To understand why the unit price of capital is pinned down at one in the growth model and to see how to break this result, it is useful to consider a more general model in which the result does not necessarily hold: a two-sector growth model with a separate investment goods production sectors. The model is the same as in section 1.4 except that the resource constraint is

$$C_t + \iota_t = \mathcal{Y}_t, \quad I_t = G_t(\iota_t), \quad K_{t+1} = I_t + (1 - \delta)K_t. \quad (44)$$

Here ι_t units of the consumption produce I_t units of investment according to a production function G_t which is increasing but which may be concave $G_t'' \leq 0$ or may vary over time. Profit maximization of investment goods producer is

$$\max_{\iota_t} p_t G_t(\iota_t) - \iota_t.$$

As long as the marginal product $G'(\iota_t)$ is positive, producers choose ι_t to satisfy the optimality condition

$$p_t G_t'(\iota_t) = 1. \quad (45)$$

This model has two interesting polar special cases:

1. Neoclassical growth model: $I_t = G(\iota_t) = \iota_t$ so that $G'(\iota_t) = 1$. In this case the resource constraints (44) become $C_t + I_t = \mathcal{Y}_t$ and the optimality condition (45) implies $p_t = 1$.
2. Capital in fixed supply (Section 6.1): $G(\iota_t) = 0$ for all ι_t and $\delta = 0$. In this case $I_t = \iota_t = 0$ so that the resource constraints (44) become $C_t = \mathcal{Y}_t$. The price of capital p_t is pinned down by market clearing $I_t = 0$ rather than the optimality condition (45) because there is no optimization problem for investment goods production.

A.1.2 Growth Model with a Stock Market

This appendix spells out in more detail a decentralization of the growth model in which households trade shares in the representative firm which are in unit fixed supply. The budget constraint of the representative household is

$$p_t(S_{t+1} - S_t) + C_t = Y_t + D_t S_t.$$

Here, each share S_t is a claim on the profits of the representative firm. In equilibrium, shares are in unit fixed supply

$$S_t = 1.$$

Denoting the wage by W_t , the firm's cash flows are

$$D_t = \mathcal{Y}_t - W_t L_t - I_t$$

Using that labor is paid its marginal product $W_t = f_L(K_t, A_t L_t) A_t$, that $f(K_t, A_t L_t) = f_K(K_t, A_t L_t) K_t + f_L(K_t, A_t L_t) A_t L_t$ because of constant returns, and $L_t = 1$:

$$D_t = f_K(K_t, A_t) K_t + (1 - \delta) K_t - K_{t+1}. \quad (46)$$

The discount rate is still given by (10) and hence $R_{t+1} = f_K(K_{t+1}, A_{t+1}) + 1 - \delta$ in equilibrium.

We next show that the share price equals the value of the capital stock $p_t = K_{t+1}$. First, as usual, the share price equals the present-discounted value of dividends, i.e. it satisfies (8) with $T = \infty$.

Lemma 4. *The share price equals the value of the capital stock*

$$p_t = K_{t+1}.$$

Proof. Because the share price satisfies (8) with $T = \infty$ it equivalently satisfies

$$p_t = R_{t+1}^{-1} (D_{t+1} + p_{t+1}).$$

Using $R_{t+1} = f_K(K_{t+1}, A_{t+1}) + 1 - \delta$ and (46)

$$p_t (f_K(K_{t+1}, A_{t+1}) + 1 - \delta) = (f_K(K_{t+1}, A_{t+1}) K_{t+1} + (1 - \delta) K_{t+1} - K_{t+2}) + p_{t+1}.$$

The p_t sequence satisfying this equation is $p_t = K_{t+1}$ as claimed. □

Remark on Lemma 4. Note that it is still true that the price per unit of capital (rather than the price of the entire capital stock $p_t = K_{t+1}$) equals one.

Balanced Growth Path (BGP). Assume that productivity A_t grows at a constant rate

$$A_{t+1} = G A_t, \quad G > 1 \quad \Rightarrow \quad A_t = G^t A_0$$

and that households have isoelastic preferences

$$U(C) = \frac{C^{1-1/\sigma}}{1-1/\sigma}.$$

Then the economy has a BGP on which the capital stock, output, and consumption all grow at the same rate G . On this BGP, the asset return is constant

$$R_{t+1} = f_K(K_{t+1}, A_{t+1}) + 1 - \delta = f_K(K_0^*, A_0) + 1 - \delta = \bar{R}$$

where we have used that $f_K(K_t, A_t)$ is homogeneous of degree zero in (K_t, A_t) . The initial location of the BGP K_0^* is pinned down by the discount rate

$$\bar{R} = \frac{1}{\beta} G^\sigma.$$

On the BGP, the asset price grows at a constant rate resulting in capital gains

$$\frac{p_{t+1}}{p_t} = \frac{K_{t+2}}{K_{t+1}} = G.$$

From (46), the dividend yield is given by

$$\frac{D_{t+1}}{p_t} = \frac{f_K(K_{t+1}, A_{t+1})K_{t+1} + (1 - \delta)K_{t+1} - K_{t+2}}{K_{t+1}} = f_K(K_{t+1}, A_{t+1}) + 1 - \delta - G = \bar{R} - G. \quad (47)$$

so that

$$\frac{D_{t+1}}{p_t} + \frac{p_{t+1}}{p_t} = \bar{R}$$

as expected. Also note that from (47) we have

$$p_t = \frac{D_{t+1}}{\bar{R} - G}$$

Therefore, all capital gains are driven entirely by growing cash flows and the price-dividend ratio is constant as in the Gordon growth model (Gordon and Shapiro, 1956). Also note that all capital gains are, in fact, unrealized. This is because, in equilibrium, the representative household does not buy or sell any shares (which are in fixed supply).

B Appendix for Section 3

B.1 Proof of Equations (20) and (21)

The optimal consumption allocation described in equations (20) and (21) for the deterministic model in Section 3 is a special case of the more general result for the stochastic economy in Section 6 (for $S = 1$), which we prove in Appendix E.4.

B.2 Proof of Proposition 1

We begin with proving the first equation in Proposition 1. From the budget constraint (22) we have

$$\begin{aligned} \bar{T}_0(\theta) &= y_0(\theta) - \bar{c}_0(\theta) + \bar{p} \bar{x}(\theta) \\ T_0(\theta) &= y_0(\theta) - c_0(\theta) + p x(\theta). \end{aligned}$$

where we denote by $\bar{c}_t(\theta)$, $t = 0, 1$, consumption at the old prices and dividends. Subtracting the former from the latter, we obtain

$$\Delta T_0(\theta) = T_0(\theta) - \bar{T}_0(\theta) = \bar{x}(\theta) \Delta p + p(x(\theta) - \bar{x}(\theta)) - (c_0(\theta) - \bar{c}_0(\theta)). \quad (48)$$

By the second-period budget constraint (23) and the normalization that $T_1(\theta)$ is held fixed, we have

$$c_1(\theta) - \bar{c}_1(\theta) = \bar{k}_1(\theta) \Delta D - D(x(\theta) - \bar{x}(\theta)) \quad (49)$$

and thus

$$p(x(\theta) - \bar{x}(\theta)) = \frac{p}{D} [\bar{k}_1(\theta)\Delta D - (c_1(\theta) - \bar{c}_1(\theta))].$$

Substituting in (48), we obtain

$$\Delta T_0(\theta) = \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D - \left[(c_0(\theta) - \bar{c}_0(\theta)) + \frac{p}{D}(c_1(\theta) - \bar{c}_1(\theta)) \right].$$

Next, recall the solution to the planning problem (19). Using $c_t(\theta) = \Omega(\theta)C_t$ we have

$$\Delta T_0(\theta) = \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D - \Omega(\theta) \left[C_0 - \bar{C}_0 + \frac{p}{D}(C_1 - \bar{C}_1) \right]. \quad (50)$$

To rewrite the expression in square brackets, we work with the aggregate resource constraints. Since $\int T_0(\theta)dF(\theta) = \int T_1(\theta)dF(\theta) = 0$ we have $C_0 = pX + Y_0$ and $C_1 = D(K_0 - X) + Y_1$. Therefore

$$\begin{aligned} C_0 - \bar{C}_0 &= pX - \bar{p}\bar{X} = \bar{X}\Delta p + p\Delta X, \\ C_1 - \bar{C}_1 &= D(K_0 - X) - \bar{D}(K_0 - \bar{X}) = K_0\Delta D - D\Delta X - \bar{X}\Delta D = \bar{K}_1\Delta D - D\Delta X \end{aligned}$$

where $\Delta X = X - \bar{X}$. Combining yields

$$C_0 - \bar{C}_0 + \frac{p}{D}(C_1 - \bar{C}_1) = \bar{X}\Delta p + \frac{p}{D}\bar{K}_1\Delta D.$$

Substituting in (50) delivers the final result:

$$\Delta T_0(\theta) = \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D - \Omega(\theta) \left[\bar{X}\Delta p + \frac{p}{D}\bar{K}_1\Delta D \right].$$

The proof of the second equation in Proposition 1 follows the same steps. Equation (48) can equivalently be written as

$$\Delta T_0(\theta) = x(\theta)\Delta p + \bar{p}(x(\theta) - \bar{x}(\theta)) - (c_0(\theta) - \bar{c}_0(\theta)) \quad (51)$$

and equation (49) as

$$c_1(\theta) - \bar{c}_1(\theta) = k_1(\theta)\Delta D - \bar{D}(x(\theta) - \bar{x}(\theta))$$

and thus

$$\bar{p}(x(\theta) - \bar{x}(\theta)) = \frac{\bar{p}}{D} [k_1(\theta)\Delta D - (c_1(\theta) - \bar{c}_1(\theta))].$$

Substituting in (51), we obtain

$$\begin{aligned} \Delta T_0(\theta) &= x(\theta)\Delta p + \frac{\bar{p}}{D}k_1(\theta)\Delta D - \left[(c_0(\theta) - \bar{c}_0(\theta)) + \frac{\bar{p}}{D}(c_1(\theta) - \bar{c}_1(\theta)) \right] \\ &= x(\theta)\Delta p + \frac{\bar{p}}{D}k_1(\theta)\Delta D - \Omega(\theta) \left[C_0 - \bar{C}_0 + \frac{\bar{p}}{D}(C_1 - \bar{C}_1) \right]. \end{aligned} \quad (52)$$

By the aggregate resource constraints

$$\begin{aligned} C_0 - \bar{C}_0 &= pX - \bar{p}\bar{X} = X\Delta p + \bar{p}\Delta X \\ C_1 - \bar{C}_1 &= D(K_0 - X) - \bar{D}(K_0 - \bar{X}) = K_0\Delta D - \bar{D}\Delta X - X\Delta D = \bar{K}_1\Delta D - \bar{D}\Delta X \end{aligned}$$

Combining yields

$$C_0 - \bar{C}_0 + \frac{\bar{p}}{D}(C_1 - \bar{C}_1) = X\Delta p + \frac{\bar{p}}{D}\bar{K}_1\Delta D.$$

Substituting in (52) delivers the final result:

$$\Delta T_0(\theta) = x(\theta)\Delta p + \frac{\bar{p}}{D}k_1(\theta)\Delta D - \Omega(\theta) \left[X\Delta p + \frac{\bar{p}}{D}K_1\Delta D \right].$$

B.3 Proof of Lemma 1

As in the Lemma, denote the old price by \bar{p} and the new price by $p = \bar{p} + \Delta p$. Similarly, denote the old dividend by \bar{D} and the new dividend by $D = \bar{D} + \Delta D$. Denote the original consumption bundle by $(\bar{c}_0(\theta), \bar{c}_1(\theta))$. Slutsky compensation is defined as the change in the investor's total budget $y_0(\theta)$ that keeps the original consumption bundle $(\bar{c}_0(\theta), \bar{c}_1(\theta))$ affordable at the new asset price p and dividend D . In the remainder of the proof, we suppress the dependence of variables on θ for notational simplicity.

The lifetime budget line at the original price is the set of points (c_0, c_1) such that

$$c_0 + \frac{\bar{p}}{D}c_1 = y_0 + \frac{\bar{p}}{D}y_1 + \bar{p}k_0 \quad (53)$$

The Slutsky-compensated budget line at the new price is the set of points (c_0, c_1) such that

$$c_0 + \frac{p}{D}c_1 = y_0 + \frac{p}{D}y_1 + pk_0 + \Delta y_0, \quad (54)$$

where Δy_0 is the Slutsky compensation term. The aim is to solve for Δy_0 such that the two budget lines intersect at the point $(c_0, c_1) = (\bar{c}_0, \bar{c}_1)$, i.e. so that the original consumption bundle remains affordable at the new prices. To this end, evaluate (53) and (54) at (\bar{c}_0, \bar{c}_1) and subtract the old budget constraint (53) from the new budget constraint (54)

$$\left(\frac{p}{D} - \frac{\bar{p}}{D} \right) \bar{c}_1 = \left(\frac{p}{D} - \frac{\bar{p}}{D} \right) y_1 + k_0\Delta p + \Delta y_0$$

Rearranging, we have

$$\Delta y_0 = \left(\frac{p}{D} - \frac{\bar{p}}{D} \right) (\bar{c}_1 - y_1) - k_0\Delta p = \left(\frac{p}{D} - \frac{\bar{p}}{D} + \frac{\Delta p}{D} \right) (\bar{c}_1 - y_1) - k_0\Delta p$$

where the second equality used $\bar{p} = p - \Delta p$. Using the second-period budget constraint (12), which implies $\bar{c}_1 - y_1 = \bar{D}\bar{k}_1$, yields

$$\Delta y_0 = \left(\frac{p}{D} - \frac{\bar{p}}{D} \right) \bar{D}\bar{k}_1 + (\bar{k}_1 - k_0)\Delta p = p \left(\frac{\bar{D}}{D} - 1 \right) \bar{k}_1 + (\bar{k}_1 - k_0)\Delta p = -\frac{p}{D}\bar{k}_1\Delta D - \bar{x}\Delta p$$

where the last equality uses the definition of asset sales in (14) which implies $\bar{x} = k_0 - \bar{k}_1$. Reintroducing the explicit dependence on θ yields the in the Lemma.

B.4 Endogenous payout policy and share repurchases

The capital-structure neutral reformulation of our setup is easiest to explain in the multi-period model of Section 1. Consider a firm that produces an income stream (i.e. earnings minus investment) $\{\Pi_t\}_{t=0}^{\infty}$ from its fundamental (e.g., non-financial) operations. Investors have budget constraint (2) and we assume for simplicity that the only asset at their disposal is firm shares so that $k_t(\theta)$ denotes share holdings, p_t denotes the share price, and D_t denotes the business dividends per share. The firm's cash flows

are distributed to shareholders through both dividends and share repurchases:

$$\Pi_t = \mathcal{K}_t D_t + (\mathcal{K}_t - \mathcal{K}_{t+1}) p_t \quad (55)$$

where and $\mathcal{K}_t = \int k_t(\theta) dF(\theta)$ denotes the total amount of outstanding shares. When $\mathcal{K}_{t+1} < \mathcal{K}_t$ the business is repurchasing its own shares. From this equation it is already apparent that share repurchases and dividend payments are equivalent means of distributing cash flows $\{\Pi_t\}_{t=0}^{\infty}$ to shareholders as a whole. When the business repurchases its shares (i.e., $\mathcal{K}_{t+1} < \mathcal{K}_t$) this results in an income stream $(k_t(\theta) - k_{t+1}(\theta)) p_t$ for those individual selling their shares to the business.

Denoting by $s_t(\theta) \equiv k_t(\theta) / \mathcal{K}_t$ the individual's ownership share of the business and by $V_t \equiv \mathcal{K}_t p_t$ the market value of the business, we can combine the individual and business budget constraints, (2) and (55), to obtain:

$$c_t(\theta) + V_t(s_{t+1}(\theta) - s_t(\theta)) = y_t(\theta) + \Pi_t s_t(\theta) \quad (56)$$

This budget constraint has the same form as (2), except that (i) the dividend per share D_t is replaced by the income stream from operations Π_t , (ii) the price per share p_t is replaced by the market value of the firm V_t , and (iii) the number of shares held by the individuals $k_t(\theta)$ is replaced by the ownership share in the business $s_t(\theta)$. An alternative viewpoint on this consolidated budget constraint is to consider the return to investing in the business. See [Fagereng et al. \(2023\)](#) for more discussion on this capital-structure neutral reformulation.

B.5 Proof of Proposition 2

In a similar manner to Proposition 1, we use the budget constraint in the first period to get

$$\begin{aligned} \bar{T}_0(\theta) &= y_0(\theta) - \bar{c}_0(\theta) - \bar{a}_1(\theta) + \bar{R}_0 a_0(\theta) \\ T_0(\theta) &= y_0(\theta) - c_0(\theta) - a_1(\theta) + R_0 a_0(\theta). \end{aligned}$$

Subtracting the former from the latter, we obtain

$$\Delta T_0(\theta) = T_0(\theta) - \bar{T}_0(\theta) = (\bar{c}_0(\theta) - c_0(\theta)) + (\bar{a}_1(\theta) - a_1(\theta)) + a_0(\theta) \Delta R_0. \quad (57)$$

From the second-period budget constraint we have

$$c_1(\theta) - \bar{c}_1(\theta) = R_1 a_1(\theta) - \bar{R}_1 \bar{a}_1(\theta). \quad (58)$$

Note that we can write it as

$$c_1(\theta) - \bar{c}_1(\theta) = \bar{a}_1(\theta)(R_1 - \bar{R}_1) + R_1(a_1(\theta) - \bar{a}_1(\theta)) \quad (59)$$

Thus we have

$$\bar{a}_1(\theta) - a_1(\theta) = \frac{1}{R_1} (\bar{c}_1(\theta) - c_1(\theta)) + \frac{1}{R_1} \bar{a}_1(\theta) \Delta R_1$$

Replacing in (57) we get

$$\Delta T_0(\theta) = a_0(\theta) \Delta R_0 + \frac{1}{R_1} \bar{a}_1(\theta) \Delta R_1 - \left[(c_0(\theta) - \bar{c}_0(\theta)) + \frac{1}{R_1} (c_1(\theta) - \bar{c}_1(\theta)) \right].$$

Again, using the solution to the planning problem (19), we know $c_t(\theta) = \Omega(\theta)C_t$. So we have

$$\Delta T_0(\theta) = a_0(\theta)\Delta R_0 + \frac{1}{R_1}\bar{a}_1(\theta)\Delta R_1 - \Omega(\theta) \left[C_0 - \bar{C}_0 + \frac{1}{R_1}(C_1 - \bar{C}_1) \right]. \quad (60)$$

Similar to Proposition 1, working with the aggregate resource constraints we have

$$C_0 + A_1 = Y_0 + R_0A_0,$$

$$C_1 = Y_1 + R_1A_1.$$

Therefore

$$C_0 - \bar{C}_0 = -\Delta A_1 + A_0\Delta R_0,$$

$$C_1 - \bar{C}_1 = R_1A_1 - \bar{R}_1\bar{A}_1 = R_1\Delta A_1 + \bar{A}_1\Delta R_1.$$

Replacing in equation (60) we get the final result

$$T_0(\theta) = \bar{T}_0(\theta) + a_0(\theta)\Delta R_0 + \frac{1}{R_1}\bar{a}_1(\theta)\Delta R_1 - \Omega(\theta) \left[A_0\Delta R_0 + \frac{1}{R_1}\bar{A}_1\Delta R_1 \right]. \quad (61)$$

The proof of the second equation in Proposition 2 follows the same steps. Equation (59) can equivalently be written as

$$c_1(\theta) - \bar{c}_1(\theta) = a_1(\theta)(R_1 - \bar{R}_1) + \bar{R}_1(a_1(\theta) - \bar{a}_1(\theta)). \quad (62)$$

Thus we have

$$\bar{a}_1(\theta) - a_1(\theta) = \frac{1}{R_1}(\bar{c}_1(\theta) - c_1(\theta)) + \frac{1}{R_1}a_1(\theta)\Delta R_1$$

Replacing in (57) we get

$$\Delta T_0(\theta) = a_0(\theta)\Delta R_0 + \frac{1}{R_1}a_1(\theta)\Delta R_1 - \left[(c_0(\theta) - \bar{c}_0(\theta)) + \frac{1}{R_1}(c_1(\theta) - \bar{c}_1(\theta)) \right].$$

Again, using the solution to the planning problem (19), we know $c_t(\theta) = \Omega(\theta)C_t$. So we have

$$\Delta T_0(\theta) = a_0(\theta)\Delta R_0 + \frac{1}{R_1}a_1(\theta)\Delta R_1 - \Omega(\theta) \left[C_0 - \bar{C}_0 + \frac{1}{R_1}(C_1 - \bar{C}_1) \right]. \quad (63)$$

Working with the aggregate resource constraints we have

$$C_0 - \bar{C}_0 = -\Delta A_1 + A_0\Delta R_0,$$

$$C_1 - \bar{C}_1 = R_1A_1 - \bar{R}_1\bar{A}_1 = \bar{R}_1\Delta A_1 + A_1\Delta R_1.$$

Replacing in equation (63) we get the final result

$$T_0(\theta) = \bar{T}_0(\theta) + a_0(\theta)\Delta R_0 + \frac{1}{R_1}a_1(\theta)\Delta R_1 - \Omega(\theta) \left[A_0\Delta R_0 + \frac{1}{R_1}A_1\Delta R_1 \right]. \quad (64)$$

B.6 Proof of Lemma 2

The proof is similar to that of Lemma 1. The lifetime budget constraint at the initial returns is

$$c_0 + \frac{c_1}{R_1} = y_0 + \frac{y_1}{R_1} + \bar{R}_0a_0. \quad (65)$$

The Slutsky-compensated budget line at the new returns is

$$c_0 + \frac{c_1}{R_1} = \Delta y_0 + y_0 + \frac{y_1}{R_1} + R_0 a_0. \quad (66)$$

As before, we evaluate (65) and (66) at (\bar{c}_0, \bar{c}_1) and subtract the old budget constraint (65) from the new budget constraint (66):

$$c_1 \left(\frac{1}{R_1} - \frac{1}{\bar{R}_1} \right) = \Delta y_0 + y_1 \left(\frac{1}{R_1} - \frac{1}{\bar{R}_1} \right) + a_0 \Delta R_0. \quad \Rightarrow \quad \Delta y_0 = \left(\frac{1}{R_1} - \frac{1}{\bar{R}_1} \right) (\bar{c}_1 - y_1) - a_0 \Delta R_0$$

Using the second period budget constraint we have

$$\Delta y_0 = \left(\frac{1}{R_1} - \frac{1}{\bar{R}_1} \right) (\bar{R}_1 \bar{a}_1) - a_0 \Delta R_0.$$

Simplifying, we get

$$\Delta y_0 = -\frac{1}{R_1} \bar{a}_1 \Delta R_1 - a_0 \Delta R_0.$$

B.7 Proof of Proposition 3

An investor's present-value budget constraint at the new prices and dividends (p, D) and new taxes $T_t(\theta)$ is

$$c_0(\theta) + \frac{p}{D} c_1(\theta) + T_0(\theta) + \frac{p}{D} T_1(\theta) = y_0(\theta) + \frac{p}{D} y_1(\theta) + p k_0(\theta) \quad (67)$$

The present-value budget constraint at the new prices and dividends (p, D) but *old* taxes $\bar{T}_t(\theta)$ is

$$\hat{c}_0(\theta) + \frac{p}{D} \hat{c}_1(\theta) + \bar{T}_0(\theta) + \frac{p}{D} \bar{T}_1(\theta) = y_0(\theta) + \frac{p}{D} y_1(\theta) + p k_0(\theta) \quad (68)$$

Subtracting (68) from (67) yields

$$c_0(\theta) - \hat{c}_0(\theta) + \frac{p}{D} (c_1(\theta) - \hat{c}_1(\theta)) + T_0(\theta) - \bar{T}_0(\theta) + \frac{p}{D} (T_1(\theta) - \bar{T}_1(\theta)) = 0$$

which we can rewrite as

$$\Delta T_0(\theta) + \frac{p}{D} \Delta T_1(\theta) = \hat{c}_0(\theta) - c_0(\theta) + \frac{p}{D} (\hat{c}_1(\theta) - c_1(\theta)) \quad (69)$$

Next, observe that, for $t = 1, 2$,

$$\hat{c}_t(\theta) - c_t(\theta) = \hat{c}_t(\theta) - \bar{c}_t(\theta) - (c_t(\theta) - \bar{c}_t(\theta)) = \Delta \hat{c}_t(\theta) - \Delta c_t(\theta) = \Delta \hat{c}_t(\theta) - \Omega(\theta) \Delta C_t,$$

where the last step uses the solution to the Pareto problem (19). Substituting back in equation (69) yields

$$\Delta T_0(\theta) + \frac{p}{D} \Delta T_1(\theta) = \Delta \hat{c}_0(\theta) - \Omega(\theta) \Delta C_0 + \frac{p}{D} (\Delta \hat{c}_1(\theta) - \Omega(\theta) \Delta C_1).$$

One way of implementing this is to set, in each period $t = 0, 1$,

$$\Delta T_t(\theta) = \Delta \hat{c}_t(\theta) - \Omega(\theta) \Delta C_t$$

as in Proposition 3.

B.8 Proof of Equation (24)

Start with the first-period budget constraint (22) holding fixed the old taxes when prices and dividends change

$$\begin{aligned}\bar{T}_0(\theta) &= y_0(\theta) - \bar{c}_0(\theta) + \bar{p}\bar{x}(\theta) \\ \bar{T}_0(\theta) &= y_0(\theta) - \hat{c}_0(\theta) + p\hat{x}(\theta)\end{aligned}$$

where we denote by $\hat{x}(\theta)$ the investor's optimal asset sales at the old taxes but new prices and dividends. Subtracting the former from the latter, we obtain

$$0 = \bar{x}(\theta)\Delta p + p(\hat{x}(\theta) - \bar{x}(\theta)) - (\hat{c}_0(\theta) - \bar{c}_0(\theta)). \quad (70)$$

By the second-period budget constraint (23) and holding the old taxes $\bar{T}_1(\theta)$ fixed, we have

$$\hat{c}_1(\theta) - \bar{c}_1(\theta) = \bar{k}_1(\theta)\Delta D - D(\hat{x}(\theta) - \bar{x}(\theta))$$

and thus

$$p(\hat{x}(\theta) - \bar{x}(\theta)) = \frac{p}{D} \left[\bar{k}_1(\theta)\Delta D - (\hat{c}_1(\theta) - \bar{c}_1(\theta)) \right].$$

Substituting in (70), we obtain

$$0 = \bar{x}(\theta)\Delta p + \frac{p}{D}\bar{k}_1(\theta)\Delta D - \left[\Delta\hat{c}_0(\theta) + \frac{p}{D}\Delta\hat{c}_1(\theta) \right].$$

C Appendix for Section 4

C.1 Second-best problem for alternative tax instruments

A tax on c_0 . Consider first a tax on period-0 consumption. This means that pre-tax consumption in period 0, given by $z_0(\theta) \equiv px(\theta) + y_0(\theta)$, is observable, so after-tax consumption is $c_0(\theta) = z_0(\theta) - T_0(z_0(\theta))$, where $T_0(z_0)$ is the nonlinear consumption tax in $t = 0$. Hence,

$$x(\theta) = \frac{z_0(\theta) - y_0(\theta)}{p}$$

and we can write the global incentive constraints as

$$\begin{aligned}U(\theta) &\equiv U\left(c_0(\theta), D\left(k_0(\theta) - \frac{z_0(\theta) - y_0(\theta)}{p}\right) + y_1(\theta)\right) \\ &\geq U\left(c_0(\hat{\theta}), D\left(k_0(\theta) - \frac{z_0(\hat{\theta}) - y_0(\theta)}{p}\right) + y_1(\theta)\right) \quad \forall \theta, \hat{\theta}.\end{aligned}$$

The local incentive constraints are therefore given by (29) with $A(\theta) = 0$ and

$$B(\theta) = Dk'_0(\theta) + \frac{D}{p}y'_0(\theta) + y'_1(\theta).$$

A tax on c_1 . Consider next a tax on period-1 consumption. Pre-tax consumption in period 1 is $z_1(\theta) \equiv D(k_0(\theta) - x(\theta)) + y_1(\theta)$ and after-tax consumption is $c_1(\theta) = z_1(\theta) - T_1(z_1(\theta))$, where $T_1(z_1)$

is the nonlinear consumption tax in $t = 1$. Hence,

$$x(\theta) = \frac{y_1(\theta) - z_1(\theta)}{D} + k_0(\theta)$$

and we can write the global incentive constraints as

$$\begin{aligned} U(\theta) &\equiv U\left(\frac{p}{D}(y_1(\theta) - z_1(\theta)) + pk_0(\theta) + y_0(\theta), c_1(\theta)\right) \\ &\geq U\left(\frac{p}{D}(y_1(\theta) - z_1(\hat{\theta})) + pk_0(\theta) + y_0(\theta), c_1(\hat{\theta})\right) \quad \forall \theta, \hat{\theta}. \end{aligned}$$

The local incentive constraints are therefore again given by (29) but with

$$A(\theta) = pk'_0(\theta) + y'_0(\theta) + \frac{p}{D}y'_1(\theta)$$

and $B(\theta) = 0$.

Further tax instruments. More generally, for any tax instrument conditioning on some observable choices, we can decompose consumption in each period $t = 0, 1$ into its observable and its unobservable components: $c_t(\theta) = c_t^o(\theta) + c_t^u(\theta)$. For instance, with an assets sales tax, the observable components are $c_0^o(\theta) = z_x(\theta)$ in period 0 and $c_1^o(\theta) = -Dx(\theta)$ in period 1, whereas the unobservable components are $c_0^u(\theta) = y_0(\theta)$ and $c_1^u(\theta) = Dk_0(\theta) + y_1(\theta)$. Hence, the general incentive constraint can be written as

$$\begin{aligned} U(\theta) &\equiv U(c_0^o(\theta) + c_0^u(\theta), c_1^o(\theta) + c_1^u(\theta)) \\ &\geq U(c_0^o(\hat{\theta}) + c_0^u(\theta), c_1^o(\hat{\theta}) + c_1^u(\theta)) \quad \forall \theta, \hat{\theta}. \end{aligned}$$

The local incentive constraints are therefore always given by (29) with $A(\theta) = c_0^{u'}(\theta)$ and $B(\theta) = c_1^{u'}(\theta)$. Note that this general approach also allows for combinations of the tax instruments considered so far. For example, suppose there is both an asset sales tax in period 0 and a wealth tax in period 1. Then $c_0^o(\theta) = z_x(\theta)$ and $c_1^o(\theta) = D(k_0(\theta) - x(\theta))$ while $c_0^u(\theta) = y_0(\theta)$ and $c_1^u(\theta) = y_1(\theta)$, so we obtain $A(\theta) = y'_0(\theta)$ and $B(\theta) = y'_1(\theta)$.³⁴

C.2 Solving the general second-best problem

For any preferences $U(c_0, c_1) = G(C(c_0, c_1))$ and any of the tax instruments considered, we can write the second-best Pareto problem as

$$\max_{\{c_0(\theta), c_1(\theta), V(\theta)\}} \int \omega(\theta) G(V(\theta)) dF(\theta)$$

subject to the incentive constraints

$$V'(\theta) = C_{c_0}(c_0(\theta), c_1(\theta))A(\theta) + C_{c_1}(c_0(\theta), c_1(\theta))B(\theta) \quad \forall \theta \quad (71)$$

where $C_{c_t} \equiv \partial C / \partial c_t$, the resource constraint

$$Y \geq \int \left(c_0(\theta) + \frac{p}{D}c_1(\theta) \right) dF(\theta) \quad (72)$$

with

$$Y \equiv pK_0 + Y_0 + \frac{p}{D}Y_1,$$

³⁴In this case, the asset sales tax and the wealth tax are not separately determined, but the optimal consumption allocation is.

and

$$V(\theta) = C(c_0(\theta), c_1(\theta)) \quad \forall \theta.$$

It is useful to substitute out $c_0(\theta) = \Phi(V(\theta), c_1(\theta))$ where $\Phi(\cdot, c_1)$ is the inverse function of $C(\cdot, c_1)$ with respect to its first argument. This allows us to write the maximization problem in terms of $V(\theta)$ and $c_1(\theta)$ only. Attaching multipliers $\mu(\theta)$ to the incentive constraint for type θ and η to the resource constraint, the corresponding Lagrangian becomes, after integrating by parts,

$$\begin{aligned} \mathcal{L} &= \int \omega(\theta) G(V(\theta)) dF(\theta) - \int \mu'(\theta) V(\theta) d\theta \\ &\quad - \int \mu(\theta) [C_0(\Phi(V(\theta), c_1(\theta)), c_1(\theta)) A(\theta) + C_1(\Phi(V(\theta), c_1(\theta)), c_1(\theta)) B(\theta)] d\theta \\ &\quad - \eta \int \left[\Phi(V(\theta), c_1(\theta)) + \frac{p}{D} c_1(\theta) \right] dF(\theta). \end{aligned}$$

Using the fact that

$$\frac{\partial \Phi}{\partial V} = \frac{1}{C_0} \quad \text{and} \quad \frac{\partial \Phi}{\partial c_1} = -\frac{U_1}{U_0}$$

and dropping arguments to simplify notation, the first-order condition for $c_1(\theta)$ is

$$\mu \left[\left(\frac{C_{c_0 c_0} C_{c_1}}{C_{c_0}} - C_{c_0 c_1} \right) A + \left(\frac{C_{c_0 c_1} C_{c_1}}{C_{c_0}} - C_{c_1 c_1} \right) B \right] = \eta f \left[\frac{p}{D} - \frac{C_{c_1}}{C_{c_0}} \right] \quad (73)$$

and for $V(\theta)$

$$\omega f G'(V) = \mu' + \mu \left[\frac{C_{c_0 c_0}}{C_{c_0}} A + \frac{C_{c_0 c_1}}{C_{c_0}} B \right] + \frac{\eta f}{C_{c_0}} \quad (74)$$

where $C_{c_s c_t}$ denotes the second derivatives $\partial^2 C / (\partial c_s \partial c_t)$. Together with the incentive constraints (71), the resource constraint (72) and the boundary conditions $\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0$, equations (73) and (74) determine the optimal solution $\{V(\theta), c_1(\theta), \mu(\theta), \eta\}$.

C.3 CES utility and numerical algorithm

Under the CES preferences given in equation (6), it turns out to be convenient to work with

$$\zeta(\theta) \equiv \frac{c_0(\theta)}{c_1(\theta)}.$$

Then the first-order conditions (73) and (74) can be written as

$$\frac{\mu(\theta)}{\sigma c_1(\theta)} \left(\zeta(\theta)^{\frac{\sigma-1}{\sigma}} + \beta \right)^{\frac{1}{\sigma-1}} (B(\theta) - A(\theta) / \zeta(\theta)) = \eta f(\theta) \left(\frac{p}{\beta D} - \zeta(\theta)^{\frac{1}{\sigma}} \right) \quad (75)$$

$$\omega(\theta) f(\theta) G'(V(\theta)) = \mu'(\theta) + \frac{\beta \mu(\theta)}{\sigma c_1(\theta)} \frac{B(\theta) - A(\theta) / \zeta(\theta)}{\zeta(\theta)^{\frac{\sigma-1}{\sigma}} + \beta} + \eta f(\theta) \left(1 + \beta \zeta(\theta)^{\frac{1-\sigma}{\sigma}} \right)^{\frac{1}{1-\sigma}}. \quad (76)$$

Moreover, the incentive constraints (71) become

$$V'(\theta) = \left(\zeta(\theta)^{\frac{\sigma-1}{\sigma}} + \beta \right)^{\frac{1}{\sigma-1}} \left(\zeta(\theta)^{-\frac{1}{\sigma}} A(\theta) + \beta B(\theta) \right) \quad (77)$$

and, by definition of CES utility,

$$V(\theta) = c_1(\theta) \left(\zeta(\theta)^{\frac{\sigma-1}{\sigma}} + \beta \right)^{\frac{\sigma}{\sigma-1}}. \quad (78)$$

We first use (75) together with (78) to numerically solve for $\xi(\theta)$ as a function of $\mu(\theta)$, $V(\theta)$ and η . Substituting this in (76) and (77) delivers a system of two ordinary differential equations in $\mu(\theta)$ and $V(\theta)$ that we can solve, given any η , using the boundary conditions $\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0$. Finally, we find η such that the resource constraint (72) is satisfied, noting that $c_0(\theta) = \xi(\theta)c_1(\theta)$.

C.4 Proof of Proposition 4

We are interested in analyzing the limiting behavior as $\sigma \rightarrow 0$ of $\{\xi(\theta), V(\theta), \mu(\theta), c_1(\theta), \eta\}$ satisfying (75) to (78), the resource constraint (72) and the boundary conditions. We make their dependence on σ explicit and write $\{\xi(\theta, \sigma), V(\theta, \sigma), \mu(\theta, \sigma), c_1(\theta, \sigma), \eta(\sigma)\}$. We continue to denote by $\mu'(\theta, \sigma)$ the derivative of μ with respect to θ .

Lemma 5. $\mu(\theta, \sigma)$ remains bounded for all θ as $\sigma \rightarrow 0$, i.e. $\mu(\theta, \sigma) \rightarrow \mu(\theta, 0)$ with $|\mu(\theta, 0)| < \infty$ for all θ .

Proof. Equations (75) and (76) simplify in the limit as $\sigma \rightarrow 0$. To see this, we use the notation $f(\sigma) \sim g(\sigma)$ for two functions f and g to indicate that $\lim_{\sigma \rightarrow 0} f(\sigma)/g(\sigma) = 1$. Then, for any $\xi(\theta) \geq 1$, as $\sigma \rightarrow 0$

$$\begin{aligned} \left(\xi(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} + \beta\right)^{\frac{1}{\sigma-1}} &= \left(\xi(\theta, \sigma)^{1-\frac{1}{\sigma}} + \beta\right)^{\frac{1}{\sigma-1}} \sim \left(\xi(\theta, \sigma)^{-\frac{1}{\sigma}} + \beta\right)^{-1} \\ \left(\xi(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} + \beta\right)^{-1} &\sim \left(\xi(\theta, \sigma)^{-\frac{1}{\sigma}} + \beta\right)^{-1} \\ \left(1 + \beta\xi(\theta, \sigma)^{\frac{1-\sigma}{\sigma}}\right)^{\frac{1}{1-\sigma}} &= \left(1 + \beta\xi(\theta, \sigma)^{\frac{1}{\sigma}-1}\right)^{\frac{1}{1-\sigma}} \sim 1 + \beta\xi(\theta, \sigma)^{\frac{1}{\sigma}}. \end{aligned}$$

Therefore (75) and (76) become

$$\frac{\mu(\theta, \sigma)}{\sigma c_1(\theta, \sigma)} \frac{B(\theta) - A(\theta)/\xi(\theta, \sigma)}{\xi(\theta, \sigma)^{-\frac{1}{\sigma}} + \beta} = \eta(\sigma) f(\theta) \left(\frac{p}{\beta D} - \xi(\theta, \sigma)^{\frac{1}{\sigma}} \right) \quad (79)$$

$$\omega(\theta) f(\theta) G'(V(\theta, \sigma)) = \mu'(\theta, \sigma) + \frac{\beta \mu(\theta, \sigma)}{\sigma c_1(\theta, \sigma)} \frac{B(\theta) - A(\theta)/\xi(\theta, \sigma)}{\xi(\theta, \sigma)^{-\frac{1}{\sigma}} + \beta} + \eta(\sigma) f(\theta) \left(1 + \beta \xi(\theta, \sigma)^{\frac{1}{\sigma}} \right) \quad (80)$$

Substituting (79) into (80) we have

$$\begin{aligned} \omega(\theta) f(\theta) G'(V(\theta, \sigma)) &= \mu'(\theta, \sigma) + \eta(\sigma) f(\theta) \left[\frac{p}{\beta D} - \xi(\theta, \sigma)^{\frac{1}{\sigma}} \right] \beta + \eta(\sigma) f(\theta) \left(1 + \beta \xi(\theta, \sigma)^{\frac{1}{\sigma}} \right) \\ &= \mu'(\theta, \sigma) + \eta(\sigma) f(\theta) \left(1 + \frac{p}{D} \right) \end{aligned}$$

Or equivalently

$$\mu'(\theta, \sigma) = \left(\omega(\theta) G'(V(\theta, \sigma)) - \eta(\sigma) \left(1 + \frac{p}{D} \right) \right) f(\theta)$$

Integrating and using the boundary condition $\mu(\underline{\theta}, \sigma) = 0$ we have

$$\mu(\theta, \sigma) = \int_{\underline{\theta}}^{\theta} \left(\omega(s) G'(V(s, \sigma)) - \eta(\sigma) \left(1 + \frac{p}{D} \right) \right) dF(s). \quad (81)$$

Because all objects on the right-hand side of this equation are bounded, we obtain Lemma 5. \square

Lemma 6. $\xi(\theta, \sigma) \rightarrow 1$ as $\sigma \rightarrow 0$.

Proof. The proof is by contradiction. Fix a θ and suppose that $\xi(\theta, \sigma)$ does not converge to 1 as $\sigma \rightarrow 0$ for that θ . There are two cases.

1. $\zeta(\theta, \sigma) \rightarrow \zeta(\theta, 0) < 1$ for that θ , i.e. $c_0(\theta, 0) < c_1(\theta, 0)$. In this case

$$\zeta(\theta, \sigma)^{\frac{1}{\sigma}} \rightarrow 0, \quad \zeta(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} = \zeta(\theta, \sigma)^{1-\frac{1}{\sigma}} \sim \zeta(\theta, \sigma)^{-\frac{1}{\sigma}} \rightarrow \infty, \quad \left(\zeta(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} + \beta\right)^{\frac{1}{\sigma-1}} \sim \zeta(\theta, \sigma)^{\frac{1}{\sigma}} \rightarrow 0$$

Using these, (75) becomes

$$\mu(\theta, \sigma) \frac{\zeta(\theta, \sigma)^{\frac{1}{\sigma}}}{\sigma} (B(\theta) - A(\theta)/\zeta(\theta, \sigma)) = \eta(\sigma) f(\theta) c_1(\theta, \sigma) \frac{p}{\beta D} \quad (82)$$

We argue that the less obvious term $\zeta(\theta, \sigma)^{\frac{1}{\sigma}}/\sigma$ on the left-hand side of (82) converges to zero as $\sigma \rightarrow 0$. Since $\zeta(\theta, \sigma) \rightarrow \zeta(\theta, 0) < 1$ as $\sigma \rightarrow 0$, there is an upper bound $\bar{\zeta}(\theta) \in (\zeta(\theta, 0), 1)$ and some $\bar{\sigma}$ such that $\zeta(\theta, \sigma) \leq \bar{\zeta}(\theta)$ for all $\sigma \leq \bar{\sigma}$. Therefore

$$0 \leq \lim_{\sigma \rightarrow 0} \frac{\zeta(\theta, \sigma)^{\frac{1}{\sigma}}}{\sigma} \leq \lim_{\sigma \rightarrow 0} \frac{\bar{\zeta}(\theta)^{\frac{1}{\sigma}}}{\sigma} \quad (83)$$

with $\bar{\zeta}(\theta) < 1$. We prove that the upper bound converges to zero as $\sigma \rightarrow 0$ using l'Hopital. To see this,

$$\lim_{\sigma \rightarrow 0} \frac{\bar{\zeta}(\theta)^{\frac{1}{\sigma}}}{\sigma} = \lim_{x \rightarrow \infty} \bar{\zeta}(\theta)^x x = \lim_{x \rightarrow \infty} \frac{x}{e^{-x \log \bar{\zeta}(\theta)}} = \lim_{x \rightarrow \infty} \frac{1}{-\log \bar{\zeta}(\theta) e^{-x \log \bar{\zeta}(\theta)}} = \lim_{x \rightarrow \infty} -\frac{\bar{\zeta}(\theta)^x}{\log \bar{\zeta}(\theta)} = 0.$$

Therefore, using (83) we also have

$$\lim_{\sigma \rightarrow 0} \frac{\zeta(\theta, \sigma)^{\frac{1}{\sigma}}}{\sigma} = 0. \quad (84)$$

From Lemma 5 we know that $\mu(\theta, \sigma)$ remains bounded as $\sigma \rightarrow 0$. Therefore taking $\sigma \rightarrow 0$ in (82) and using (84) yields

$$0 = \eta(0) f(\theta) c_1(\theta, 0) \frac{p}{\beta D}$$

which is a contradiction because $c_1(\theta, 0) > 0$ and $f(\theta) > 0$ for at least some θ and $\eta(0) > 0$. We have thus ruled out case 1.

2. $\zeta(\theta, \sigma) \rightarrow \zeta(\theta, 0) > 1$ for that θ , i.e. $c_0(\theta, 0) > c_1(\theta, 0)$. In this case

$$\zeta(\theta, \sigma)^{\frac{1}{\sigma}} \rightarrow \infty, \quad \zeta(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} = \zeta(\theta, \sigma)^{1-\frac{1}{\sigma}} \sim \zeta(\theta, \sigma)^{-\frac{1}{\sigma}} \rightarrow 0, \quad \left(\zeta(\theta, \sigma)^{\frac{\sigma-1}{\sigma}} + \beta\right)^{\frac{1}{\sigma-1}} \rightarrow \beta^{-1}$$

Using this, we can write equation (75) as

$$\frac{\mu(\theta, \sigma) \zeta(\theta, \sigma)^{-\frac{1}{\sigma}}}{\beta \sigma} (B(\theta) - A(\theta)/\zeta(\theta, \sigma)) = \eta(\sigma) c_1(\theta, \sigma) f(\theta) \left(\frac{p}{\beta D} \zeta(\theta, \sigma)^{-\frac{1}{\sigma}} - 1 \right). \quad (85)$$

We argue that

$$\frac{\zeta(\theta, \sigma)^{-\frac{1}{\sigma}}}{\sigma} \rightarrow 0.$$

Since $\zeta(\theta, \sigma) \rightarrow \zeta(\theta, 0) > 1$ as $\sigma \rightarrow 0$, there is a lower bound $\underline{\zeta}(\theta) \in (1, \zeta(\theta, 0))$ and some $\bar{\sigma}$ such that $\zeta(\theta, \sigma) \geq \underline{\zeta}(\theta)$ for all $\sigma \leq \bar{\sigma}$. Therefore

$$0 \leq \lim_{\sigma \rightarrow 0} \frac{\zeta(\theta, \sigma)^{-\frac{1}{\sigma}}}{\sigma} \leq \lim_{\sigma \rightarrow 0} \frac{\underline{\zeta}(\theta)^{-\frac{1}{\sigma}}}{\sigma} \quad (86)$$

with $\underline{\zeta}(\theta) > 1$. Finally, we prove that the upper bound in (86) converges to zero using l'Hopital:

$$\lim_{\sigma \rightarrow 0} \frac{\underline{\zeta}(\theta)^{-\frac{1}{\sigma}}}{\sigma} = \lim_{x \rightarrow \infty} \underline{\zeta}(\theta)^{-x} x = \lim_{x \rightarrow \infty} \frac{x}{e^{x \log \underline{\zeta}(\theta)}} = \lim_{x \rightarrow \infty} \frac{1}{\log \underline{\zeta}(\theta) e^{x \log \underline{\zeta}(\theta)}} = \lim_{x \rightarrow \infty} \frac{\underline{\zeta}(\theta)^{-x}}{\log \underline{\zeta}(\theta)} = 0$$

Therefore, using (86) we also have

$$\lim_{\sigma \rightarrow 0} \frac{\underline{\zeta}(\theta, \sigma)^{-\frac{1}{\sigma}}}{\sigma} = 0. \quad (87)$$

Moreover, from Lemma 5 we know that $\mu(\theta, \sigma)$ remains bounded as $\sigma \rightarrow 0$. Hence, (85) becomes

$$0 = -\eta(0)f(\theta)c_1(\theta, 0)$$

which is a contradiction because $c_1(\theta, 0) > 0$ and $f(\theta) > 0$ for at least some θ , and $\eta(0) > 0$. We have thus ruled out case 2.

Since both cases lead to a contradiction we conclude that instead $\zeta(\theta, \sigma) \rightarrow 1$ as $\sigma \rightarrow 0$. \square

[To be completed]

D Appendix for Section 5

D.1 Proof of Proposition 5

An investor's sequential budget constraint under the old price and dividend path $\{\bar{p}_t\}_{t=0}^T$ and $\{\bar{D}_t\}_{t=0}^T$ is

$$\bar{T}_t(\theta) = y_t(\theta) + \bar{D}_t \bar{k}_t(\theta) + \bar{p}_t(\bar{k}_t(\theta) - \bar{k}_{t+1}(\theta)) - \bar{c}_t(\theta)$$

and under the new price and dividend path $\{p_t\}_{t=0}^T$ and $\{D_t\}_{t=0}^T$:

$$T_t(\theta) = y_t(\theta) + D_t k_t(\theta) + p_t(k_t(\theta) - k_{t+1}(\theta)) - c_t(\theta)$$

Subtracting one from the other yields:

$$\Delta T_t(\theta) = k_t(\theta) \Delta D_t + x_t(\theta) \Delta p_t + (\bar{D}_t + \bar{p}_t) \Delta k_t(\theta) - \bar{p}_t \Delta k_{t+1}(\theta) - \Delta c_t(\theta)$$

where $x_t(\theta) = k_t(\theta) - k_{t+1}(\theta)$. Premultiplying by $\bar{R}_{0 \rightarrow t}^{-1}$ and rearranging yields

$$\bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) = \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] + \bar{R}_{0 \rightarrow t}^{-1} (\bar{D}_t + \bar{p}_t) \Delta k_t(\theta) - \bar{R}_{0 \rightarrow t}^{-1} \bar{p}_t \Delta k_{t+1}(\theta)$$

Note that

$$\bar{R}_{0 \rightarrow t}^{-1} (\bar{D}_t + \bar{p}_t) = \frac{\bar{D}_t + \bar{p}_t}{\bar{R}_{0 \rightarrow t}} = \frac{\bar{D}_t + \bar{p}_t}{\bar{R}_{0 \rightarrow t-1} \frac{\bar{D}_t + \bar{p}_t}{\bar{p}_{t-1}}} = \frac{\bar{p}_{t-1}}{\bar{R}_{0 \rightarrow t-1}}$$

Therefore

$$\bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) = \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] + \bar{R}_{0 \rightarrow t-1}^{-1} \bar{p}_{t-1} \Delta k_t(\theta) - \bar{R}_{0 \rightarrow t}^{-1} \bar{p}_t \Delta k_{t+1}(\theta)$$

Summing from $t = 0$ to $t = T$, we obtain

$$\begin{aligned} \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) &= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] \\ &\quad + \sum_{t=0}^T \bar{R}_{0 \rightarrow t-1}^{-1} \bar{p}_{t-1} \Delta k_t(\theta) - \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \bar{p}_t \Delta k_{t+1}(\theta) \end{aligned}$$

Using $\Delta k_0(\theta) = 0$ and changing the index in the second sum on the right-hand side,

$$\begin{aligned} \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) &= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] \\ &\quad + \sum_{t=0}^{T-1} \bar{R}_{0 \rightarrow t}^{-1} \bar{p}_t \Delta k_{t+1}(\theta) - \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \bar{p}_t \Delta k_{t+1}(\theta) \end{aligned}$$

Therefore

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] - \bar{R}_{0 \rightarrow T}^{-1} \bar{p}_T \Delta k_{T+1}(\theta)$$

Since $k_{T+1}(\theta) = 0$, the last term vanishes. Therefore,

$$\begin{aligned} \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) &= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - \Delta c_t(\theta)] \\ &= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t] - \Omega(\theta) \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta C_t \end{aligned} \quad (88)$$

where the second line uses the optimal consumption allocation.

Subtracting the sequential resource constraints under the new and old price and dividend path yields

$$\Delta C_t = X_t \Delta p_t + K_t \Delta D_t + (\bar{D}_t + \bar{p}_t) \Delta K_t - \bar{p}_t \Delta K_{t+1}$$

Following the same steps as before yields

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta C_t = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [X_t \Delta p_t + K_t \Delta D_t]$$

Substituting back in (88), we obtain Proposition 5.

D.2 Proof of Corollary 3

Under condition (31), we have

$$\begin{aligned}
\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} \Delta T_t(\theta) &= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [k_t(\theta)(\Delta p_t + \Delta D_t) - k_{t+1}(\theta)\Delta p_t - \Omega(\theta)(K_t(\Delta p_t + \Delta D_t) - K_{t+1}\Delta p_t)] \\
&= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [k_t(\theta) - \Omega(\theta)K_t] (\Delta p_t + \Delta D_t) \\
&\quad - \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [k_{t+1}(\theta) - \Omega(\theta)K_{t+1}] \frac{\bar{p}_t}{\bar{D}_{t+1} + \bar{p}_{t+1}} (\Delta p_{t+1} + \Delta D_{t+1}) \\
&= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [k_t(\theta) - \Omega(\theta)K_t] (\Delta p_t + \Delta D_t) \\
&\quad - \sum_{t=0}^T \bar{R}_{0 \rightarrow t+1}^{-1} [k_{t+1}(\theta) - \Omega(\theta)K_{t+1}] (\Delta p_{t+1} + \Delta D_{t+1}) \\
&= \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [k_t(\theta) - \Omega(\theta)K_t] (\Delta p_t + \Delta D_t) - \sum_{t=1}^{T+1} \bar{R}_{0 \rightarrow t}^{-1} [k_t(\theta) - \Omega(\theta)K_t] (\Delta p_t + \Delta D_t) \\
&= [k_0(\theta) - \Omega(\theta)K_0] (\Delta p_0 + \Delta D_0) - \bar{R}_{0 \rightarrow T+1}^{-1} [k_{T+1}(\theta) - \Omega(\theta)K_{T+1}] (\Delta p_{T+1} + \Delta D_{T+1}) \\
&= [k_0(\theta) - \Omega(\theta)K_0] \Delta p_0
\end{aligned}$$

since the last term vanishes and $\Delta D_0 = 0$.

D.3 Multi-period version of Proposition 2 (taxing total capital income)

Proposition 9. *Suppose asset prices change by $\{\Delta p_t\}_{t=0}^T$ and dividends by $\{\Delta D_t\}_{t=0}^T$ resulting in return changes $\{\Delta R_t\}_{t=0}^T$. Then optimal taxes $\{T_t(\theta)\}_{t=0}^T$ are such that*

$$\sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} T_t(\theta) = \sum_{t=0}^T \bar{R}_{0 \rightarrow t}^{-1} [\bar{T}_t(\theta) + a_t(\theta)\Delta R_t - \Omega(\theta)A_t\Delta R_t]$$

Proof. The proof follows exactly analogous steps to the proof of Proposition 5. □

E Appendix for Section 6

E.1 Proof of Equation (34)

Recall the Euler equation (3) under preferences (6):

$$c_0(\theta)^{-1/\sigma} = \beta R c_1(\theta)^{-1/\sigma}$$

with $R = D/p$. Since it holds for all investors, it aggregates to

$$C_1 = \left(\beta \frac{D}{p} \right)^\sigma C_0.$$

Moreover, integrating the budget constraints (11) and (12) across investors and using the market clearing condition (33), we obtain $C_0 = Y_0$ and $C_1 = Y_1 + DK$ in the closed economy. Substituting back in

the aggregate Euler equation, the equilibrium asset price p^* must satisfy

$$Y_1 + DK = \left(\beta \frac{D}{p^*} \right)^\sigma Y_0,$$

which can be rearranged to deliver equation (34).

E.2 Proof of Proposition 6

First observe that, given preferences (6), the optimal consumption allocation still satisfies $c_t(\theta) = \Omega(\theta)C_t$, $t = 0, 1$. Since $C_0 = Y_0$ and $C_1 = Y_1 + DK$ in the closed economy and we hold both dividends and the aggregate endowment fixed, this immediately implies that no investor's consumption is changing in response to the asset price change Δp^* , so $c_t(\theta) = \bar{c}_t(\theta)$ for all θ , $t = 0, 1$. By the second-period budget constraint (23) and the normalization $T_1(\theta) = 0$, this implies in turn that $x(\theta) = \bar{x}(\theta)$ for all θ . The result then follows from Proposition 1 and the fact that $\Delta D = 0$ and $X = K_0 - K_1 = 0$.

E.3 Proof of Proposition 7

Subtract an investor's budget constraints under the old and new prices in period 0:

$$\begin{aligned} c_0(\theta) - \bar{c}_0(\theta) + q(b(\theta) - \bar{b}(\theta)) &= px(\theta) - \bar{p}\bar{x}(\theta) - (\chi(x(\theta)) - \chi(\bar{x}(\theta))) - (T_0(\theta) - \bar{T}_0(\theta)) \\ &= (p - \bar{p})x(\theta) + \bar{p}(x(\theta) - \bar{x}(\theta)) - (\chi(x(\theta)) - \chi(\bar{x}(\theta))) - (T_0(\theta) - \bar{T}_0(\theta)) \end{aligned}$$

and in period 1:

$$c_1(\theta) - \bar{c}_1(\theta) = D(\theta)(x(\theta) - \bar{x}(\theta)) + b(\theta) - \bar{b}(\theta)$$

We eliminate $b(\theta) - \bar{b}(\theta)$ by substituting the latter into the former:

$$\begin{aligned} c_0(\theta) - \bar{c}_0(\theta) + q(c_1(\theta) - \bar{c}_1(\theta)) + qD(\theta)(x(\theta) - \bar{x}(\theta)) \\ = (p - \bar{p})x(\theta) + \bar{p}(x(\theta) - \bar{x}(\theta)) - (\chi(x(\theta)) - \chi(\bar{x}(\theta))) - (T_0(\theta) - \bar{T}_0(\theta)) \end{aligned}$$

Rearranging and using (35) yields

$$\begin{aligned} c_0(\theta) - \bar{c}_0(\theta) + q(c_1(\theta) - \bar{c}_1(\theta)) - \chi'(\bar{x}(\theta))(x(\theta) - \bar{x}(\theta)) \\ = (p - \bar{p})x(\theta) - (\chi(x(\theta)) - \chi(\bar{x}(\theta))) - (T_0(\theta) - \bar{T}_0(\theta)) \end{aligned}$$

The second-order Taylor approximation for $\chi(x)$ around the point $\bar{x}(\theta)$ is:

$$\chi(x(\theta)) - \chi(\bar{x}(\theta)) \approx \chi'(\bar{x}(\theta))(x(\theta) - \bar{x}(\theta)) + \frac{1}{2}\chi''(\bar{x}(\theta))(x(\theta) - \bar{x}(\theta))^2$$

Substituting this, we obtain

$$c_0(\theta) - \bar{c}_0(\theta) + q(c_1(\theta) - \bar{c}_1(\theta)) = x(\theta)\Delta p - \frac{1}{2}\chi''(\bar{x}(\theta))(\Delta x(\theta))^2 - (T_0(\theta) - \bar{T}_0(\theta)) \quad (89)$$

where $\Delta x(\theta) \equiv x(\theta) - \bar{x}(\theta)$.

Since the aggregate resource constraint (37) takes the same form as (17), the Pareto problem (19)

subject to (37) still implies $c_t(\theta) = \Omega(\theta)C_t$, $t = 0, 1$, under preferences (6). Hence, (89) can be written as

$$\Delta T_0(\theta) = x(\theta)\Delta p - \frac{1}{2}\chi''(\bar{x}(\theta))(\Delta x(\theta))^2 - \Omega(\theta) [C_0 - \bar{C}_0 + q(C_1 - \bar{C}_1)]$$

Integrating (89) across all investors implies

$$C_0 - \bar{C}_0 + q(C_1 - \bar{C}_1) = X\Delta p - \frac{1}{2} \int \chi''(\bar{x}(\theta))\Delta x(\theta)^2 dF(\theta)$$

and substituting this back delivers Proposition 7.

E.4 Proof of Equation (41)

We characterize the solution to the planning problem (40). Letting $\eta(s)$ denote the multiplier on the resource constraint for state $s \in S$, the first-order conditions for $c_0(\theta)$ and $c_1(\theta, s)$ are:

$$\begin{aligned} \omega(\theta)U_0 &= \sum_{s'} \eta(s') \\ \omega(\theta)U_1 \frac{\partial \mu(\cdot)}{\partial c_1(\theta, s)} &= q\eta(s), \quad \forall s \in S, \end{aligned}$$

where $U_0 = \partial U / \partial c_0$ and $U_1 = \partial U / \partial \mu$. This implies that

$$\frac{U_0}{U_1 \partial \mu / \partial c_1(\theta, s)} = \frac{\sum_{s'} \eta(s')}{q\eta(s)}$$

is independent of θ . Given the functional forms, we have that $c_1^*(\theta, s) / c_0^*(\theta)$ is independent of θ . That is, the ratio of consumption between periods is equalized across individuals and equal to the aggregate ratio $C_1^*(s) / C_0^*$. In particular, $c_1(\theta, s) = \zeta(s)c_0(\theta)$ for some constant of proportionality $\zeta(s)$ that is common across agents. Substituting into the certainty-equivalent aggregator μ , we have $\mu(\{c_1^*(\theta, s)\}) = c_0^*(\theta)\mu(\{\zeta(s)\})$, and hence the first-order condition for c_0 becomes

$$\omega(\theta)c_0(\theta)^{-\gamma} = \chi,$$

where χ is independent of θ . Integrating over θ we obtain

$$C_0 = \chi^{-1/\gamma} \int \omega(\theta)^{\frac{1}{\gamma}} dF(\theta),$$

and

$$c_0^*(\theta) = \Omega(\theta)C_0^*,$$

where $\Omega(\theta) = \omega(\theta)^{\frac{1}{\gamma}} / \int \omega(\theta)^{\frac{1}{\gamma}} dF(\theta)$. Given that $c_1^*(\theta, s) / c_0(\theta) = C_1(s) / C_0$, we have that $c_1^*(\theta, s) = \Omega(\theta)C_1^*(s)$, as well. Finally, given that the ratio of first and second period consumptions are equalized across θ , it is equivalent to maximize first over the aggregates C_0 and $C_1(s)$ and then allocate across θ according to the optimal rule.

E.5 Proof of Lemma 3

In the decentralized equilibrium, the individual's problem is

$$\begin{aligned} \max U(c_0(\theta), \{c_1(\theta, s)\}) \quad \text{s.t.} \\ c_0(\theta) &= y_0(\theta) + px(\theta) - qb(\theta) - T_0(\theta) \\ c_1(\theta, s) &= b + D(s)(k_0(\theta) - x(\theta)) - T_1(\theta, s) \quad \forall s \in S. \end{aligned}$$

Eliminating b , we can write the budget set as S present-value constraints (suppressing θ):

$$c_0 + qc_1 = y_0 + qy_1 + px + qD(s)(k_0 - x) - T_0 - T_1(s) \quad \forall s \in S.$$

The first-order conditions for this problem take the same form as the planning problem, so the first-best allocation satisfies the individual's problem as long as it satisfies the budget set. For this, we need to find a tax scheme $\{T_0, T_1(\cdot, s)\}$ and asset positions $\{b, x\}$ such that for all θ and s :

$$\begin{aligned} \Omega(\theta)C_0^* &= y_0(\theta) + px(\theta) - qb(\theta) - T_0(\theta) \\ \Omega(\theta)C_1^*(s) &= y_1(\theta) + D(s)(k_0(\theta) - x(\theta)) + b(\theta) - T_1(\theta, s). \end{aligned}$$

Using $T_1(\theta, s) = a(\theta) + m(\theta)D(s)$ and the aggregate resource constraint in state s , we have:

$$\Omega(\theta) (D(s)K_1^* + Y_1 + B^*) = y_1(\theta) + D(s)(k_0(\theta) - x(\theta)) + b(\theta) - a(\theta) - m(\theta)D(s).$$

For each $s \in S$, set

$$\begin{aligned} x(\theta) &= -m(\theta) - \Omega(\theta)K_1^* - k_0(\theta) \\ b(\theta) &= \Omega(\theta)(Y_1 + B^*) - y_1(\theta) + a(\theta), \end{aligned}$$

and the second-period budget constraint is satisfied for all $s \in S$. Setting

$$T_0(\theta) = \Omega(\theta)C_0^* + y_0(\theta) + px(\theta) - qb(\theta),$$

the first-period budget constraint is satisfied, as well. Moreover, we have $\int T_0(\theta)dF = 0$. Hence, the tax system along with the proposed policies $\{b, x\}$ ensures that the household's necessary and sufficient conditions are satisfied evaluated at the first-best allocation.

E.6 Proof of Proposition 8

Recall that given aggregates $\{C_0, \{C_1(s)\}\}$, the first-best allocation can be written:

$$\begin{aligned} c_0(\theta) &= \Omega(\theta)C_0 \\ c_1(\theta, s) &= \Omega(\theta)C_1(s) \quad \forall s, \end{aligned}$$

where $\Omega(\theta)$ is given by (21). Thus, the change in the first-best allocation is

$$\begin{aligned} \Delta c_0(\theta) &= \Omega(\theta)\Delta C_0 \\ \Delta c_1(\theta, s) &= \Omega(\theta)\Delta C_1(s). \end{aligned}$$

From the aggregate resource condition:

$$\begin{aligned}\Delta C_0 &= pX - \bar{p}\bar{X} - qB + \bar{q}\bar{B} \\ &= X\Delta p + \bar{p}\Delta X - B\Delta q - \bar{q}\Delta B,\end{aligned}$$

and

$$\begin{aligned}\Delta C_1(s) &= D(s)(K_0 - X) + B - \bar{D}(s)(K_0 - \bar{X}) - \bar{B} \\ &= (K_0 - X)\Delta D(s) + \bar{D}(s)\Delta X + \Delta B.\end{aligned}$$

We now show that the first-best allocation is affordable for each θ . Let individual θ alter first-period saving by:

$$\Delta b(\theta) = \Omega(\theta)\Delta B,$$

and capital sales:

$$\Delta x(\theta) = \Omega(\theta)\Delta X.$$

If the first-best allocation is attainable with these portfolio choices and the proposed lump-sum transfers, they are consistent with individual optimization. This follows from the fact that the first-best allocation optimizes the distribution of consumption across time and states for each θ given the budget set.

The proposed tax for period zero is:

$$\Delta T_0(\theta) = x(\theta)\Delta p - b(\theta)\Delta q - \Omega(\theta)(X\Delta p - B\Delta q).$$

Then, under the proposed tax policy:

$$\begin{aligned}\Delta c_0(\theta) &= px(\theta) - \bar{p}\bar{x}(\theta) - qb + \bar{q}\bar{b} - \Delta T_0(\theta) \\ &= x(\theta)\Delta p + \bar{p}\Delta x(\theta) - b(\theta)\Delta q - \bar{q}\Delta b(\theta) - \Delta T_0(\theta) \\ &= x(\theta)\Delta p - b(\theta)\Delta q + \bar{p}\Omega(\theta)\Delta X - \bar{q}\Omega(\theta)\Delta B \dots \\ &\quad - [x(\theta)\Delta p - b(\theta)\Delta q - \Omega(\theta)(X\Delta p - B\Delta q)] \\ &= \Omega(\theta)(X\Delta p - B\Delta q - \bar{q}\Delta B + \bar{p}\Delta X) \\ &= \Omega(\theta)\Delta C_0.\end{aligned}$$

For the second period in state s , the proposed tax is:

$$\Delta T_1(\theta, s) = (k_0(\theta) - x(\theta))\Delta D(s) - \Omega(\theta)((K_0 - X)\Delta D(s)).$$

The change in the budget constraint in the second period becomes:

$$\begin{aligned}
\Delta c_1(\theta, s) &= D(s)(k_0(\theta) - x(\theta)) - \bar{D}(s)(k_0(\theta) - \bar{x}(\theta)) + \Delta b - \Delta T_1(\theta, s) \\
&= (k_0(\theta) - x(\theta))\Delta D(s) - \bar{D}(s)\Delta x(\theta) + \Delta b - \Delta T_1(\theta, s) \\
&= (k_0(\theta) - x(\theta))\Delta D(s) - \bar{D}(s)\Omega(\theta)\Delta X + \Omega(\theta)\Delta B \\
&\quad - [(k_0(\theta) - x(\theta))\Delta D(s) - \Omega(\theta)((K_0 - X)\Delta D(s))] \\
&= \Omega(\theta) ((K_0 - X)\Delta D(s) - \bar{D}(s)\Delta X + \Delta B) \\
&= \Omega(\theta)C_1(\theta, s).
\end{aligned}$$

Hence, the proposed taxes allow each θ to afford the change in the first-best allocation.

E.7 Proof of Corollary 5

$\Delta M(s) = 0, \forall s \in S$ implies that

$$\begin{aligned}
\Delta p &= \sum_s \pi(s)\bar{M}(s)\Delta D(s) \\
\Delta q &= 0.
\end{aligned}$$

Substituting into Proposition 8, we have

$$\begin{aligned}
T_0(\theta) &= \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p \\
T_1(\theta, s) &= \bar{T}_1(\theta, s) + (k_0(\theta) - x(\theta))\Delta D(s) - \Omega(\theta)(K_0 - X)\Delta D(s).
\end{aligned}$$

We can exploit Lemma 3 and consider an alternative tax system, as stated in the corollary:

$$\begin{aligned}
\hat{T}_0(\theta) &= \bar{T}_0(\theta) + k_0(\theta)\Delta p - \Omega(\theta)K_0\Delta p \\
\hat{T}_1(\theta, s) &= \bar{T}_1(\theta, s).
\end{aligned}$$

In the latter, relative to the former, agent θ pays $(k_1(\theta) - \Omega(\theta)K_1)\Delta p$ more in taxes in the first-period. Instead, she pays $(k_1(\theta) - \Omega(\theta)K_1)\Delta D(s)$ less in taxes in period 1 in every state s . Reducing its investment in risky capital by the amount of extra taxes in period 0, the agent perfectly undoes the tax changes in both periods and each state of the world. Hence, the budget sets are invariant to this transformation, and either allows the agent to afford the first-best allocation.