

PUTTING THE ‘FINANCE’ INTO ‘PUBLIC FINANCE’:
A THEORY OF CAPITAL GAINS TAXATION

MARK AGUIAR

Department of Economics, Princeton University

BENJAMIN MOLL

Department of Economics, London School of Economics

FLORIAN SCHEUER

Department of Economics, University of Zurich

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 movements, 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 a combination of realization-based capital gains and cash flow taxes implements the optimal allocation regardless of the source of asset-price fluctuations. Moreover, the capital gains tax avoids distortions in portfolio choice (the so-called lock-in effect) by targeting total net trades rather than gains from selling individual assets. These results stand 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 Alan Auerbach, V.V. Chari, Antoine Ferey, Mike Golosov, Camille Landais, Dmitry Mukhin, Chris Phelan, Jim Poterba, Daniel Reck, Karl Schulz, Dejanir Silva, Joel Slemrod, Juan Carlos Suarez Serrato, Aleh Tsyvinski, Iván Werning, Danny Yagan, and seminar participants at ANU, Bank of Italy, Bank of Portugal, Berkeley, Berlin, BI Oslo, Carnegie Mellon, Chicago, Columbia, EEA, IFS, LMU Munich, LSE, Mannheim, MIT, NBER, Northwestern, NYU, Rotterdam, SED, Stanford, St. Gallen, TSE, Venice, Yale and Zurich 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.

“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 never conclusively ruled on whether unrealized gains constitute income.² Because wealth changes due to asset-price movements typically dwarf ordinary saving and income flows for top wealth holders, debates about wealth taxation also often end up being about the desirability (and practicality) of taxing unrealized capital gains.

The existing public finance literature on optimal capital taxation abstracts from explicitly modelling asset prices, and therefore provides no guidance in these debates.³ Our paper aims to fill this gap 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

We show that optimal redistributive taxes generally differ from the case with constant asset prices. While there are many tax implementations of the optimal allocation, there always exists a particularly simple and robust one that targets realized trades rather than asset holdings. This

¹U.S. Office of Management and Budget (2022), U.S. Department of the Treasury (2022), [Saez et al. \(2021\)](#), and [Zucman \(2024\)](#). [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.

²This is despite the Supreme Court having repeatedly heard such cases since *Eisner v. Macomber* in 1920. The key 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, [Atkinson and Stiglitz \(1976\)](#), [Chamley \(1986\)](#), and [Judd \(1985\)](#). Like us, [Piketty et al. \(2023\)](#) lament the absence of asset-price effects from the literature. 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.

⁴See for example [Shiller \(1981\)](#), [Campbell and Shiller \(1988\)](#), [Cochrane \(2011\)](#), [Greenwald et al. \(2019\)](#) and [van Binsbergen \(2020\)](#), 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.

implementation, which is a combination of realization-based capital gains and dividend taxes, applies without requiring knowledge of the source of asset-price changes. In contrast, taxing unrealized capital gains is suboptimal whenever asset-price changes are not exclusively driven by cash flow changes. The intuition 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 takes 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 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

Our study of redistributive taxation with changing asset prices starts from a simple baseline environment: a small open economy in which a large number of investors trade two assets (risky capital and a risk-free bond) with exogenously given asset prices and asset returns that are homogeneous across investors. Investors begin with heterogeneous endowments of the two assets 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, we study various extensions, including general equilibrium and heterogeneous returns, and show that the results from this simple setting generalize.

We are interested in how the optimal tax system redistributes in response to changing 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 an investor's tax burden should react 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

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

⁵After 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.

where p_t denotes the asset’s price and D_t its cash flow, i.e. the return equals dividend yield plus capital gain. Suppose the economy is initially in a steady state with a constant asset price \bar{p} , dividend \bar{D} , and associated asset return \bar{R} . This is an example of the case typically studied in the literature and the properties of optimal capital tax systems in such steady states are well understood (see the literature discussion below). We instead allow $\{p_t, D_t, R_t\}$ to fluctuate in flexible ways. Suppose that at time $t = 0$ a shock hits the economy and results in asset prices, returns and cash flows deviating from the initial steady state. For example, asset prices p_t may increase because expected future cash flows D_t increase or for other reasons that are independent of changes in cash flows D_t , i.e. discount rate changes. The question we are after is: how should the tax system redistribute in response 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. Dividend income changes are similarly compensated or taxed. Building on the Slutsky-compensation logic, optimal first-best taxation is straightforward: just like Slutsky compensation, it taxes sellers, compensates buyers, and taxes dividend income changes. Importantly, it generally targets 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} . 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} , corresponding to asset prices driven entirely by cash flows.

We show that, in the first special case with changing discount rates, the change in the tax burden depends *only* on investors’ realized trades (purchases and sales) and the price changes relative to steady state—it is 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. A simple implementation that works in both special (and all intermediate) cases is a realization-based capital gains tax combined with a dividend tax.

Wealth taxes are sometimes likened to taxes on “presumptive income” (Zucman, 2024, or the Dutch “box 3” wealth tax): for example, a 2% wealth tax is equivalent to a 40% tax on *presumed* capital income from a *constant* asset return of 5%. When asset values increase and the increase is entirely due to higher cashflows (the second special case), the asset return remains constant and therefore the increase in presumptive income exactly matches the increase in

actual income. But in all other cases, the return falls and therefore actual income rises by less than presumptive income calculated as a constant return to the increased market value of wealth. Thus “presumptive income” is overestimated and hence wealth taxes redistribute suboptimally whenever asset valuations are not exclusively driven by cash flows.⁶

One asset class where the disconnect between income and asset valuation can be especially pronounced is startups. These firms often generate little or no current income but may nevertheless command staggering valuations (e.g., “unicorn” valuations of \$1bn). Furthermore, these valuations then frequently collapse to zero down the road – see, e.g., [Azevedo et al. \(2025\)](#) who document that founders receive zero exit value after more than 80% of U.S. venture capital deals. Taxes on wealth or unrealized capital gains would tax startup owners on these paper valuations. By contrast, our theory implies that taxes should only be levied on cashflows and in case of a successful exit.

While our formula for optimal redistributive taxes is reminiscent of realization-based capital gains taxes in practice, it also differs 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 incurs no tax liability); (iv) they adjust for inflation; and (v) 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 the first-best case with lump-sum taxes, our formula corresponds to a tax rate of 100%, i.e. the government taxes away realized capital gains in their entirety and uses the proceeds to compensate the losers from rising prices.

The first-best tax scheme is designed for redistribution, not to replace missing insurance markets. As a result, there exist other implementations, including one that achieves all redistribution in the initial period by ensuring that all investors hold the market portfolio and are therefore equally affected by future asset price or cash flow fluctuations. Our tax scheme focusing on net trades, however, does have important advantages in terms of simplicity and robustness over ones that involve transfers only in the initial period, a point we discuss in more detail later on.

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 investor *choices* can be taxed, for example asset sales, consumption, or savings. Our interest remains how the second-best optimal policy redistributes when asset prices change. We show that the 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 preferred distribution of income, like in the first-best, but now also distorts saving behavior. We also show that the optimal tax schedule converges to the first-best optimum as the investors’ inter-temporal elasticity of substitution goes to zero. Hence, our insights from the first-best

⁶Supplementary Appendix F further illustrates this analogy by means of a simple numerical example.

tax system are not knife-edge, but extend qualitatively to environments with more limited and realistic tax instruments.

If 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.

Our analysis of second-best tax systems also considers the “lock-in” effect emphasized in the capital gains taxation literature: realization-based taxes may incentivize deferring the liquidation of appreciated assets and distort optimal portfolio allocation. Using a two-asset version of our model, we show that an optimally designed second-best tax system avoids such distortions even when it targets realized capital gains. It does so by targeting total *net* trades rather than gains from selling individual assets: when an investor sells one asset and uses the proceeds to purchase another one, there is no tax burden, thus eliminating the lock-in effect.

Finally, we consider several extensions: general equilibrium, return heterogeneity, and bequests. While some of these features modify our optimal tax formula in natural ways, the key findings emphasized so far remain unchanged. Specifically, in contrast to the classic Haig-Simons comprehensive income tax concept, there generally exists an implementation of optimal redistributive taxes that targets realized trades. In fact, we show that when investors receive heterogeneous cash flows, taxes on unrealized capital gains or wealth are no longer optimal *even when asset prices are driven entirely by cash flows*, reinforcing our results from the baseline setting. Moreover, in general equilibrium, it will typically be the case that changes in cash flows also change the stochastic discount factor, further undermining the Haig-Simons approach to taxation. Our model with borrowing also speaks 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” strategy. Our results suggest that basis step-up should be abolished, eliminating the viability of such plans.

Literature. Our paper contributes to the literature studying the optimal taxation of capital income and wealth.⁷ To differentiate our paper, 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) or works with variants of the neoclassical growth model. In this 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 that are independent of changes in cash flows, i.e. discount rate

⁷Apart from the classic contributions mentioned above, see the references in Section 1.5 and the surveys by Golosov et al. (2007), Banks and Diamond (2010), Bastani and Waldenstrom (2020), Stantcheva (2020), and Scheuer and Slemrod (2021).

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. Instead, in all such cases, a combination of realization-based capital gains and dividend taxes still implements the optimal allocation.

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 (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., forthcoming](#), [Coven et al., 2024](#)). The logic of our results is related to [Moll \(2020\)](#) and [Fagereng et al. \(forthcoming\)](#) 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.⁸ 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, but such elasticities will be key inputs in more quantitative work.

While the modern capital taxation literature provides no guidance on how to tax capital gains, an older literature anticipates some of the ideas in our paper using verbal or graphical arguments. This includes [Paish \(1940\)](#), [Kaldor \(1955\)](#) and [Whalley \(1979\)](#), which were partly reactions to [Haig \(1921\)](#) and [Simons \(1938\)](#) who developed the eponymous income concept.

Roadmap. Section 1 spells out our baseline environment. Section 2 focuses on a special case with two time periods and no risk to convey our key results most transparently. Section 3 studies the first-best allocation assuming that the government has access to type-specific lump-sum taxes. In contrast, Section 4 considers the second-best problem with distortive taxation and discusses the lock-in effect. Section 5 shows how our findings carry over to the stochastic multi-period model of Section 1. Section 6 considers extensions and Section 7 concludes.

1. BASELINE MODEL

We begin by spelling out our baseline environment, which is kept purposely simple: A large number of investors trade two assets (risky capital and a risk-free bond) in a small open economy with exogenous asset prices and returns that are homogeneous across investors. In Section 6, we will consider various extensions, including general equilibrium, return heterogeneity,

⁸See for example [Poterba \(2002\)](#), [Feldstein et al. \(1980\)](#), [Agersnap and Zidar \(2021\)](#), and [Msall and Næss \(2025\)](#). 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](#)).

and intergenerational considerations. For now, we omit taxes from the analysis, which we will introduce in Section 3.

1.1. Investors

Time is discrete and indexed by $t = 0, 1, \dots, T$, where T may be finite or infinite. Let s_t denote the state of nature in period t , which takes discrete values in a set \mathcal{S} , and let $s^t \equiv \{s_0, s_1, \dots, s_t\}$ denote the history of states up to and including period t with associated probabilities $\pi(s^t)$. When convenient, we shall suppress the history notation and simply use the time index.

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 sequences, $\{c_t(\theta, s^t)\}_{t,s^t}$, captured by the utility function $U(\{c_t(\theta, s^t)\})$, which is assumed to be homothetic, strictly increasing, strictly concave and differentiable in all its arguments. We embed the probability of history s^t inside the function U , which allows us to nest both expected utility and other popular specifications, such as [Epstein and Zin \(1989\)](#). Investors receive type-specific exogenous income flows $\{y_t(\theta)\}_{t=0}^T$. Since we focus on wealth and capital gains, we assume for simplicity that the income paths are deterministic for each type.

Households can transfer income across periods by saving in two assets: a potentially risky asset k that pays a dividend stream $\{D_t(s^t)\}_{t,s^t}$ and a risk-free, zero-coupon bond b . For now, we also take prices as given, with $p_t(s^t)$ denoting the price of capital and $q_t(s^t)$ the price of the bond in period t . Investors are endowed with initial assets $\{k_0(\theta, s_{-1}), b_0(\theta, s_{-1})\}$ at time zero and, at each history s^t , choose a portfolio $\{k_{t+1}(\theta, s^t), b_{t+1}(\theta, s^t)\}$ to carry into the next period. There is no short-selling constraint, and hence asset positions may be positive or negative. In particular, investors can have negative bond holdings (i.e., borrow) while at the same time owning the capital asset. In Section 5.2 we will use this setup to discuss optimal taxation when investors borrow against appreciating assets.

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

$$\begin{aligned} \mathcal{U}(\theta) = & \max_{\{c_t(\theta, s^t), k_{t+1}(\theta, s^t), b_{t+1}(\theta, s^t)\}_{t,s^t}} U(\{c_t(\theta, s^t)\}_{t,s^t}) \quad \text{subject to} \\ & c_t(\theta, s^t) + p_t(s^t)(k_{t+1}(\theta, s^t) - k_t(\theta, s^{t-1})) + q_t(s^t)b_{t+1}(\theta, s^t) \\ & = y_t(\theta) + D_t(s^t)k_t(\theta, s^{t-1}) + b_t(\theta, s^{t-1}) \quad \forall t, s^t. \end{aligned} \quad (2)$$

We impose $p_T(s^T) = q_T(s^T) = 0$ if T is finite or a No-Ponzi condition if $T = \infty$.

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

$$C_t(s^t) = \int c_t(s^t, \theta) dF(\theta), \quad K_t(s^{t-1}) = \int k_t(s^{t-1}, \theta) dF(\theta), \quad B_t(s^{t-1}) = \int b_t(s^{t-1}, \theta) dF(\theta),$$

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

$$p_t(s^t)K_{t+1}(s^t) + q_t(s^t)B_{t+1}(s^t) + C_t(s^t) = (p_t(s^t) + D_t(s^t))K_t(s^{t-1}) + B_t(s^{t-1}) + Y_t \quad (3)$$

for all t, s^t . As already noted, our benchmark analysis focuses on a small open economy with an exogenously given time path for asset prices and dividends $\{q_t(s^t), p_t(s^t), D_t(s^t)\}_{t, s^t}$. Hence, the economy's aggregate bond and capital holdings at time $t + 1$ may differ from those at t as the economy as a whole may be a net buyer or net seller of B or K . In Section 6.1, we alternatively consider a closed-economy general equilibrium version of the model in which the assets are in fixed supply, so that 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 emphasizes 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 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.⁹ Other studies have argued that fluctuations in cash flows are first order.¹⁰ Our reading of this debate is that it is imperative to understand the tax implications of both sources.

Our partial equilibrium model takes dividends $\{D_t(s^t)\}$ and asset prices $\{p_t(s^t), q_t(s^t)\}$ as given. Instead, the perspective of the asset pricing literature is to treat required asset returns or stochastic discount rates as a primitive and prices as an outcome. One way of thinking about this is that, in equilibrium models, it is typically the discount factor that is pinned down which, in turn, determines asset prices. To this end, retain the small-open-economy assumption and denote by $m_{t+1}(s^{t+1})$ the stochastic discount factor of the representative counterparty in global financial markets between history s^t and s^{t+1} . To simplify notation, we drop the s^t -arguments and write m_{t+1} . Using this, the asset prices in period t satisfy

$$q_t = \mathbb{E}_t [m_{t+1}], \quad (4)$$

$$p_t = \mathbb{E}_t [m_{t+1} (D_{t+1} + p_{t+1})]. \quad (5)$$

⁹See [Campbell \(2018, Section 5.3.1\)](#) for an expository derivation of the Campbell-Shiller decomposition.

¹⁰See for example, [Larrain and Yogo \(2008\)](#) and [Atkeson et al. \(2024\)](#), but also see [Nagel \(2024\)](#).

In words, since a bond purchased in period t pays off one unit of consumption in all states of the world in period $t + 1$, its price is given by the mean stochastic discount factor. Similarly, the price of the risky asset at time t equals the expected discounted sum of dividend and price at $t + 1$, i.e., it consists of dividend yield and capital gain. Dividing both sides by p_t and using the definition of the asset return R_{t+1} in equation (1) yields $1 = \mathbb{E}_t[m_{t+1}R_{t+1}]$, so the stochastic discount factor m_{t+1} and R_{t+1} are inversely related. Defining

$$m_{t \rightarrow t+k} \equiv m_{t+1} \cdot m_{t+2} \cdots m_{t+k}$$

as the stochastic discount factor between history s^t and s^{t+k} and iterating on equation (5) yields

$$p_t = \mathbb{E}_t \left[\sum_{k=1}^{T-t} m_{t \rightarrow t+k} D_{t+k} \right], \quad (6)$$

i.e., the asset price equals the expected present-discounted value of future dividends. The asset price may therefore change for two reasons: changing dividends $\{D_{t+k}\}$ or changing discount factors $\{m_{t \rightarrow t+k}\}$. Accordingly, we can consider the following two extremal cases:

1. Changes in asset prices $\{p_t, q_t\}$ driven entirely by changes in the stochastic discount factor $\{m_t\}$ while holding dividends $\{D_t\}$ fixed. An important special case occurs when discount rates change such that the bond price $\{q_t\}$ and hence the risk-free interest rate $\{1/q_t\}$ remain unchanged, which corresponds to a pure risk-premium change.
2. Changes in the asset prices $\{p_t\}$ driven entirely by changes in dividends $\{D_t\}$ while holding the stochastic discount factor $\{m_t\}$ fixed (the bond prices $\{q_t\}$ remain constant in this case). An important special case is the Gordon growth model (Gordon and Shapiro, 1956) or stochastic versions of it.

Both of these cases are the opposite extremes of the general, intermediate case, with arbitrary changes in $\{p_t, q_t\}$ and $\{D_t\}$, which corresponds to asset price changes driven by a mixture of dividend and discount rate changes.

Thus, our approach is flexible enough to capture a wide range of state-of-the-art asset pricing models. For example, since by equations (4) and (5) the asset prices $\{p_t, q_t\}$ depend on the probabilities $\pi(s^t)$ with which different histories occur (through the expectations operator), we can also allow for subjective beliefs as potential drivers of asset prices (Adam et al., 2017, Bordo et al., 2023). More optimistic beliefs correspond to putting higher subjective probabilities $\tilde{\pi}(s^t)$ on histories in which cash flows $D_t(s^t)$ are high. This generates an increase in the asset price p_t while leaving actual cash flows unaffected. Changes in subjective beliefs are therefore equivalent to discount rate shocks (Special Case 1).

1.4. The deterministic case

While our results hold for the general model we have just introduced, some insights become particularly easy to understand in the special case without uncertainty ($|S| = 1$). Then (4) and (5) imply $q_t = m_{t+1} = 1/R_{t+1}$, so the bond and the capital asset are equivalent and the model collapses to the single-asset case. Furthermore, rather than considering different realizations of random variables, the deterministic case lends itself to simple comparative statics exercises which can be interpreted as realizations of MIT shocks.

Comparative statics as MIT shocks. In the deterministic case, equation (6) simplifies to

$$p_t = \sum_{k=1}^{T-t} \frac{D_{t+k}}{R_{t \rightarrow t+k}}, \quad (7)$$

where $R_{t \rightarrow t+k}$ is the cumulative return between time t and $t+k$:

$$R_{t \rightarrow t+k} \equiv R_{t+1} \cdot R_{t+2} \cdots R_{t+k}. \quad (8)$$

Below we will often conduct comparative statics in which the time path of some variable changes from a baseline to an alternative, which then induces a change in asset prices. For example, in Special Case 2 above, dividends may change from $\{\bar{D}_t\}$ to $\{D_t\} = \{\bar{D}_t + \Delta D_t\}$ holding constant $\{R_t\} = \{\bar{R}_t\}$, resulting in a change in the time path of asset prices

$$\Delta p_t = \sum_{k=1}^{T-t} \bar{R}_{t \rightarrow t+k}^{-1} \Delta D_{t+k}. \quad (9)$$

Alternatively, in Special Case 1 above, asset prices may change without any corresponding change in dividends. One useful interpretation of such comparative statics is as MIT shocks, i.e. realizations of zero-probability events in a stochastic setting: at time t , some new information arrives which changes asset prices going forward. As we show when we analyze the fully stochastic model in Section 5, our expressions take the same form regardless of whether we conduct comparative statics in a deterministic setting or compare across histories in a stochastic setting. For this reason, our analysis in Sections 2 to 4 focuses on the deterministic case.

Haig-Simons income. As already noted, in the deterministic case, the model collapses to the single-asset case. Dropping the bond, we can write the budget constraint (2) as

$$c_t(\theta) + p_t(k_{t+1}(\theta) - k_t(\theta)) = y_t(\theta) + D_t k_t(\theta) \quad \forall t \geq 0, \quad (10)$$

which 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):

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

Formulation (10) features disposable income, whereas this formulation features “Haig-Simons income” which includes unrealized capital gains (Haig, 1921, Simons, 1938). Defining the market value of wealth $a_t(\theta) \equiv p_{t-1} k_t(\theta)$ and the net return including capital gains $r_t \equiv R_t - 1$, Haig-Simons income also equals $y_t(\theta) + r_t a_t(\theta)$, i.e. income including total capital income.¹¹ Similarly, adding $a_t(\theta)$ on both sides of the budget constraint yields the standard

$$c_t(\theta) + a_{t+1}(\theta) = y_t(\theta) + R_t a_t(\theta) \quad \forall t \geq 0, \quad (11)$$

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

1.5. 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 $\{p_t, D_t\}$, and hence returns $\{R_t\}$.

Partial equilibrium models. This is the special case with $R_t = \bar{R}$ for all t . The most obvious way of generating this 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).

Neoclassical growth model. Starting with Chamley (1986), many papers have studied optimal capital taxation in variants of the growth model. Denote by $\sum_{t=0}^T \beta^t U(C_t)$ 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 any case, 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):

$$R_{t+1} = \frac{1}{\beta} \frac{U'(C_t)}{U'(C_{t+1})}. \quad (12)$$

¹¹The 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 $a_t(\theta) \equiv p_t k_t(\theta)$.

Furthermore, the unit price of capital (relative to consumption) equals one because the consumption good can be converted into investment one-for-one.¹²

In contrast, dividends and asset prices $\{D_t, p_t\}$ differ across decentralizations. For example, our asset may correspond to 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](#)).¹³ 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 Supplementary Appendix A so variations in the capital stock generate 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 income is the sum of dividend income plus (unrealized) capital gains and Chamley's result is that the long-run tax rate on this combined capital income should be zero.

Our interest is in optimal capital gains taxation in response to asset-price fluctuations away from such a balanced growth path. With the right decentralization, a stochastic version of the model above (as in, for example, [Zhu, 1992](#), [Chari and Kehoe, 1999](#)) would feature such asset-price fluctuations. However, with standard shock processes, movements in the stochastic discount factor and hence asset return R_{t+1} would be quantitatively small, analogous to the disappointing asset-pricing properties of the real business cycle model.¹⁴ We instead allow for flexible stochastic processes for the drivers of asset prices, including potentially large fluctuations in stochastic discount factors and asset returns.

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 Werning \(2020\)](#), [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. In sum, the setups studied in the existing literature feature either constant asset prices or small movements of asset returns and a constant unit price of capital. We instead

¹²We discuss this property in more detail in Supplementary Appendix A where we also discuss how to break it.

¹³[Chari et al. \(2018\)](#), an earlier version of [Chari et al. \(2020\)](#), analyzes optimal capital taxation in such a setup.

¹⁴The difficulty with explaining asset returns in RBC models is 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.”

study optimal taxation with exogenous stochastic processes for discount factors and dividends $\{m_t, D_t\}$ and associated prices and returns $\{p_t, R_t\}$. This 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), Sections 5 and 6 show that our findings remain valid in richer environments endogenizing these fluctuations.

1.6. Efficient allocations

We conclude this section by evoking an auxiliary property of first-best Pareto efficient allocations in our general model that will be useful below. For now we do not consider the question of implementing these allocations with taxes, which will be the focus of the next sections.

Let $\omega(\theta)$ be the Pareto weight on an investor of type θ . Any Pareto efficient allocation $\{c_t^*(s^t, \theta)\}$ must satisfy the following sub-problem:

$$\max_{\{c(s^t, \theta)\}} \int \omega(\theta) U(\{c(s^t, \theta)\}) dF(\theta) \quad \text{s.t.} \quad \int c(s^t, \theta) dF(\theta) \leq C_t(s^t) \quad \forall t, s^t \quad (13)$$

for some aggregate consumption $\{C_t(s^t)\}$. The following lemma shows that, given our assumption of homothetic preferences, any optimum will obey a linear sharing rule.

LEMMA 1: *Any solution $\{c_t^*(s^t, \theta)\}$ to problem (13) satisfies*

$$c_t^*(s^t, \theta) = \Omega(\theta) C_t(s^t) \quad \forall \theta, t, s^t$$

for some θ -dependent constant $\Omega(\theta)$ that is increasing in $\omega(\theta)$ and satisfies $\int \Omega(\theta) dF(\theta) = 1$.

This property will allow for particularly transparent expressions in our benchmark results below on how first-best optimal taxes respond to changes in asset prices. In Section 4, we will consider second-best optimal taxes where this property no longer necessarily holds.

2. TWO TIME PERIODS

To build intuition, we now turn to the case with two time periods and no risk. Our analysis of optimal taxation in the next two sections uses this model before we show in Section 5 that our findings carry over to the general multi-period case with uncertainty.

2.1. The investor’s problem

With two time periods $t = 0, 1$ and no risk, we can drop the bond (as discussed in Section 1) and the investor’s problem is to maximize utility $U(c_0(\theta), c_1(\theta))$ subject to:

$$c_0(\theta) + p(k_1(\theta) - k_0(\theta)) = y_0(\theta), \quad (14)$$




FIGURE 1.—Example of heterogeneity in initial assets and incomes over time

$$c_1(\theta) = y_1(\theta) + Dk_1(\theta). \quad (15)$$

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

$$R_0 \equiv \frac{p}{p_{-1}} \quad \text{and} \quad R_1 \equiv \frac{D}{p}, \quad (16)$$

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

A useful reformulation of the investors' problem is in terms of asset sales $x(\theta) \equiv k_0(\theta) - k_1(\theta)$, where $x > 0$ represents sales and $x < 0$ purchases of k . Using this, the investors solve:

$$\begin{aligned} \mathcal{U}(\theta) &\equiv \max_{\{c_0(\theta), c_1(\theta), x(\theta)\}} U(c_0(\theta), c_1(\theta)) \quad \text{s.t.} \\ c_0(\theta) &= y_0(\theta) + px(\theta) \\ c_1(\theta) &= y_1(\theta) + D(k_0(\theta) - x(\theta)) \end{aligned} \quad (17)$$

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, $k_0(\theta) - x(\theta)$.

Fundamentally, investors differ in initial asset holdings $k_0(\theta)$ and incomes $\{y_0(\theta), y_1(\theta)\}$. This heterogeneity generates gains from trade, with natural buyers and sellers of the asset. Figure 1 depicts an example 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 are buyers of the asset with $x(\theta) < 0$ whereas high- θ types are sellers with $x(\theta) > 0$. Effectively, low- θ types, who have a more front-loaded income profile, are savers

(with $c_0(\theta) < y_0(\theta)$) whereas high- θ types, who have more future income, are borrowers (with $c_0(\theta) > y_0(\theta)$).¹⁵

It is sometimes useful to combine the two period budget constraints in (17) into a present-value budget constraint:

$$c_0(\theta) + \frac{p}{D}c_1(\theta) = y_0(\theta) + \frac{p}{D}y_1(\theta) + pk_0(\theta). \quad (18)$$

This constraint states that the present-discounted value of consumption (discounted at the asset return $R_1 = D/p$ defined in (16)) equals the present-discounted value of income plus initial wealth. This constraint can also be aggregated across all investors to yield

$$C_0 + \frac{p}{D}C_1 \leq Y_0 + \frac{p}{D}Y_1 + pK_0, \quad (19)$$

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

Sources of asset-price changes. In this model, the discussion of cash flows and discount rates as drivers of asset-price changes in Section 1.3 becomes particularly simple. Treating the required return R_1 and dividends D as the primitives in (16), the two-period version of (7) is simply $p = D/R_1$. Special Case 1 is thus the case in which the asset price p changes holding dividends D fixed. On the opposite extreme, Special Case 2 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 how the optimal tax system deals with changing 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 [Dávila and Korinek \(2018\)](#), [Moll \(2020\)](#) and [Fagereng et al. \(forthcoming\)](#), we use the envelope theorem to differentiate the value function $\mathcal{U}(\theta)$ of investors defined in (17) to obtain

$$d\mathcal{U}(\theta) = U_{c_0}(\theta) \left(x(\theta) dp + \frac{p}{D}k_1(\theta) dD \right). \quad (20)$$

Consider first Special Case 1: the asset price rises ($dp > 0$) but cash flows are fixed ($dD = 0$). The marginal welfare effect 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 sellers of the asset (i.e., $x(\theta) > 0$) and hurts buyers (i.e., $x(\theta) < 0$). To first order, it does not affect individuals who do not plan to trade (i.e., $x(\theta) = 0$):

¹⁵For example, when interpreting the asset as a bond, (2) implies the budget constraints $c_0(\theta) = y_0(\theta) + b_0(\theta) - qb_1(\theta)$ and $c_1(\theta) = y_1(\theta) + b_1(\theta)$ where q is the bond price. Thus, those with a steeper income profile (e.g., more future human capital) are borrowers who sell bonds ($b_1(\theta) < 0$) and vice versa.

for them, the increasing asset price is merely a “paper gain” with no corresponding effect on welfare. Hence, only asset *transactions* matter whereas asset *holdings* do not.

When dividends rise ($dD > 0$)—as in Special Case 2—the second term in (20) becomes non-zero since this directly benefits asset holders ($k_1(\theta) > 0$). However, it remains true that the 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 redistributes in response to changing asset prices. As a first step, we will assume that the government has access to type-specific lump-sum taxes. While this implies extreme predictions about tax *rates*, it turns out to be instructive about the optimal tax *base*, i.e., what quantities taxes should target, which is our main object of interest. We will consider more realistic, second-best tax systems in Section 4.

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

$$\max_{\{c_0(\theta), c_1(\theta)\}} \int \omega(\theta) U(c_0(\theta), c_1(\theta)) dF(\theta) \quad \text{s.t.} \quad (19). \quad (21)$$

By Lemma 1, we have $c_t^*(\theta) = \Omega(\theta)C_t^*$, so the planner assigns to each investor θ a time-invariant share of (optimally chosen) aggregate consumption C_t^* . 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 (14) and (15) 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 is to set $T_1(\theta) = 0$. Then the second-period budget constraint (23) determines $x^*(\theta)$ and we obtain $T_0(\theta)$ from the first-period budget constraint (22). However, we will also consider alternative normalizations below when this is convenient.

3.2. Taxing changing asset prices

We begin with the general case of prices and dividends (p, D) that vary relative to some baseline values (\bar{p}, \bar{D}) , allowing for asset prices driven by a mixture of discount rate and dividend changes. The goal is to design a tax rule $T_0(\theta) = T_0(\theta; p, D)$ that optimally redistributes across investors in response to these (p, D) variations. We denote by $\bar{T}_0(\theta) = T_0(\theta; \bar{p}, \bar{D})$ the taxes that implement the Pareto efficient allocation at the baseline prices and dividends, and by $\Delta p = p - \bar{p}$ and $\Delta D = D - \bar{D}$ the changes in prices and dividends relative to the baseline.

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 $T_1(\theta)$ is held fixed) 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 $x(\theta)$ and $k_1(\theta)$ are investor θ 's asset sales and second-period asset holdings at the new price p and dividends D , X and K_1 are the corresponding aggregate asset sales and holdings, and, similarly, $\bar{x}(\theta)$, $\bar{k}_1(\theta)$, \bar{X} , and \bar{K}_1 are asset sales and holdings at the baseline price \bar{p} and dividend \bar{D} .

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 2: *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

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 (20) 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. 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.

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 shows that the response of the optimal tax system to an asset price change is closely related to the Slutsky compensation from Lemma 2. Indeed, the two would coincide if there were no aggregate asset trade with the rest of the world, $\bar{X} = 0$. The 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 welfare weights $\Omega(\theta)$.

According to Lemma 2, 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 $-D/p$). 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.

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 precisely 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 under homothetic preferences.

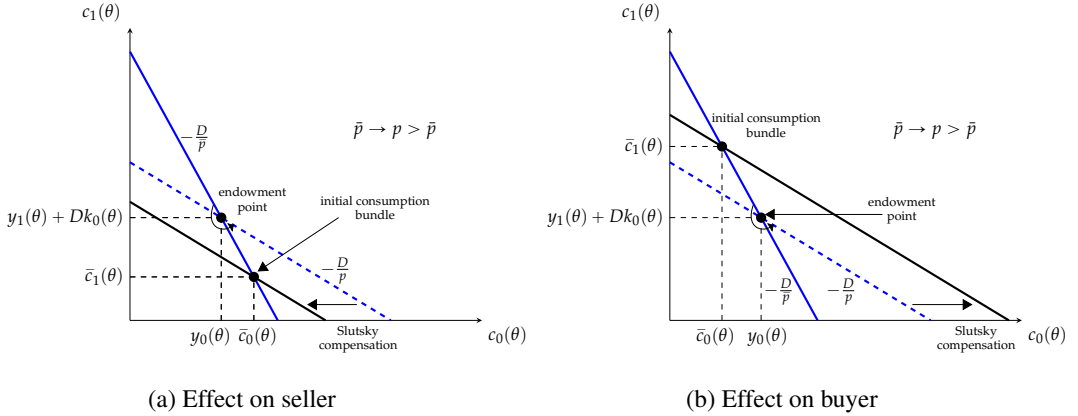


FIGURE 2.—Slutsky compensation after a pure asset-price increase (Special Case 1)

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 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, $x(\theta) = \bar{x}(\theta)$.¹⁸

When $X = 0$, as happens under some parameter configurations¹⁹ 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.²⁰ 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, optimal taxes target “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

¹⁸This happens in the closed economy of Section 6.1 in which aggregate asset sales are zero. In this case, optimal policy simply takes everyone back to their baseline consumption allocation, which implies $x(\theta) = \bar{x}(\theta)$.

¹⁹For example, with preferences $u(c_0) + \beta u(c_1)$, $u' > 0, u'' < 0$, $X = 0$ if $\beta \bar{D}/p = 1$ and $Y_0 = Y_1 + \bar{D}K_0$.

²⁰Our theory features only real variables, so it prescribes taxing inflation-indexed capital gains (as practiced in some countries like Israel).

tax conditions on *net* transactions only. For instance, if an individual sells a house and then buys a similar house of the same price, and house prices go up, she realizes a capital gain on the sale 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. We will return to this in Section 4.²¹

General Case

We now return to the general case in Proposition 1. In addition to the terms discussed so far, 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 is complemented by a dividend income tax.

This is again closely related to the Slutsky compensation in Lemma 2. 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—see Figure 3a.

While the intuition is therefore similar to the welfare gains formula (20), 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 (20)) play no role. Since budget constraints are linear in prices, Lemma 2 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

$$\begin{aligned} T_0(\theta) &= \bar{T}_0(\theta) + \bar{x}(\theta)\Delta p - \Omega(\theta)\bar{X}\Delta p = \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p \\ T_1(\theta) &= \bar{T}_1(\theta) + \bar{k}_1(\theta)\Delta D - \Omega(\theta)\bar{K}_1\Delta D = \bar{T}_1(\theta) + k_1(\theta)\Delta D - \Omega(\theta)K_1\Delta D, \end{aligned}$$

²¹As mentioned in Section 2, selling the asset corresponds to borrowing and buying the asset to saving. Thus, when interpreting the asset as a bond, the analogue of Corollary 1 is that, when the bond price changes by Δq , the optimal change in the tax $T_0(\theta)$ is $\Delta T_0(\theta) = -b_1(\theta)\Delta q + \Omega(\theta)B_1\Delta q$. In other words, when the discount rate falls ($\Delta q > 0$), those who borrow ($b_1(\theta) < 0$), for instance, because they have more future human capital (i.e., a steeper income profile), should see their tax burden increase ($\Delta T_0(\theta) > 0$), and vice versa. We will discuss this in more detail in the dynamic, multi-asset model in Section 5.

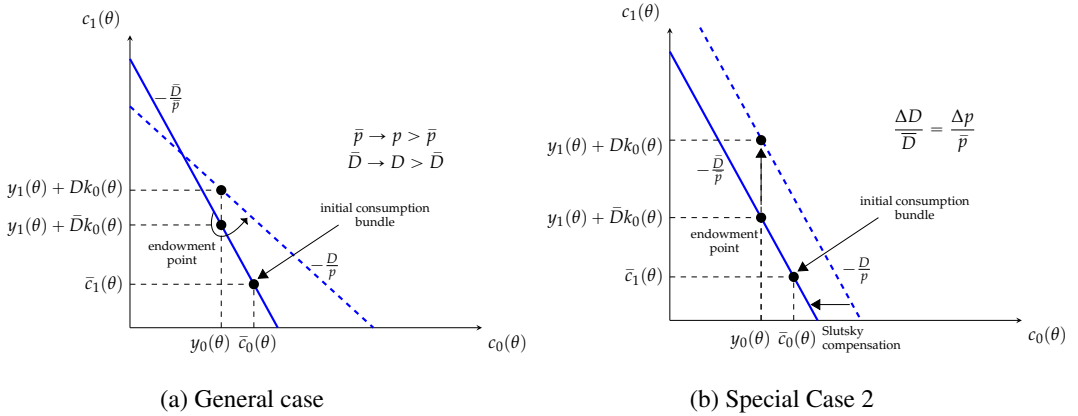


FIGURE 3.—Slutsky compensation of combined asset-price and cash-flow changes

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.

where $\{\bar{T}_t(\theta)\}$, $t = 0, 1$, are some optimal baseline taxes. 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.

Special Case 2: Only Cash Flow Changes

Finally, we consider the second extreme case from Section 1.3 in which asset prices change exclusive because of 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 individuals sells. Formally, we can use the first equation in Proposition 1 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.

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 $(y_0(\theta), y_1(\theta) + Dk_0(\theta))$ shifts upward. In the lifetime budget constraint (18), 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 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. While the “trading gains and losses” $x(\theta)\Delta p$ bear similarities to realized capital gains (in case of an asset sale), an important difference is that the price change Δp is relative to some baseline price \bar{p} , which 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.

In the general model with uncertainty from Section 1, a natural baseline price and dividend would be given by the corresponding means. Hence, the old allocation can be interpreted as the optimum under these expected prices and dividends whereas 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 baseline is a Gordon

growth model or BGP in which the asset return is constant and dividends grow at a constant rate (see Section 1).²²

Baseline taxes. Proposition 1 also assumes that taxes are set optimally for given Pareto weights at the baseline prices and dividends. If baseline taxes were not set optimally (or based on different Pareto weights), we could always decompose the overall change in taxes into two steps: First, holding baseline prices and dividends fixed, a reform of the baseline taxes towards the optimum according to the new Pareto weights. Second, holding Pareto weights fixed, a move towards the optimum under the new prices and dividends. 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. Supplementary Appendix B provides an alternative, capital-structure neutral formulation of our setup in which such distinctions are immaterial. 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 flow of housing services and implementing our tax formula therefore requires valuing this “dividend.” The solution is to measure 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 such as Denmark and Switzerland.

3.4. *Alternative implementations: taxes on expenditure or total capital income*

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 consumption and another one based on total capital income.

²²Observed capital gains taxes instead use historical purchase prices to compute capital gains upon sale. Imposing this institutional feature would yield a third-best, Ramsey-style optimal tax problem rather than the clean compensation scheme we study here. We return to this point in Section 5.

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, 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 an expenditure-based tax. Supplementary Appendix B formalizes this conjecture. It shows that the new optimum after a change in asset price and dividends can be implemented with a combination of lump-sum taxes and transfers targeting consumption changes. Notably, if the parameter changes Δp and ΔD are “zero-sum,” so that optimal aggregate consumption C_t does not change, then optimal redistributive taxation simply taxes away any increase in consumption from the asset-price and dividend changes (or compensates the corresponding reduction in consumption), i.e. a “pure” expenditure tax. In line with Kaldor’s logic, just like Proposition 1, this works for any combination of asset price and dividend changes, i.e. regardless of the source of capital gains.

In Supplementary Appendix B, we also show that yet another way of writing the first-best tax response in Proposition 1 is in terms of investors’ market value of wealth $a_t(\theta) \equiv p_{t-1}k_t(\theta)$ and the changes in the total returns R_0 and R_1 : In each period, the additional total capital income $a_t(\theta)\Delta R_t$, including unrealized gains, is taxed. While this may, at first glance, appear related to a Haig-Simons notion of income, we discuss in Supplementary Appendix B that there are important differences. For instance, a one-off permanent increase in the asset price increases the return in period 0 (leading to a tax), but reduces the return in period 1 due to the reduced dividend-price ratio (leading to a subsidy). This can lead to very volatile taxes compared to Proposition 1.

4. SECOND BEST

We now turn to the case where the government’s tax instruments are more limited. Specifically, they are restricted to condition on investors’ choices, such as their asset sales, wealth, or consumption. This distorts investors’ behavior, inducing the classic tradeoff between redistribution and efficiency. Our main conclusion is that the previous results on the tax base generalize in a natural way. We first consider a one-asset setup as above, isolating savings distortions. Sec-

tion 4.4 then considers multiple assets and the question whether taxes create a “lock-in” effect that distorts portfolio choice.

4.1. Mirrlees Problem

An asset sales tax. We begin with a (non-linear) tax $T_x(px)$ on asset sales, paid in period 0, similar to a realization-based capital gains tax. The investors’ budget constraints become

$$c_0(\theta) = y_0(\theta) + px(\theta) - T_x(px(\theta)) \quad \text{and} \quad c_1(\theta) = y_1(\theta) + D(k_0(\theta) - x(\theta)).$$

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

$$\mathcal{U}(\theta) \equiv U(z_x(\theta) + y_0(\theta), D(k_0(\theta) - x(\theta)) + y_1(\theta)) \geq U(z_x(\hat{\theta}) + y_0(\theta), D(k_0(\theta) - x(\hat{\theta})) + y_1(\theta))$$

for all $\theta, \hat{\theta}$. Abstracting from bunching, we work with the local version of the incentive constraints. By the envelope theorem,

$$\mathcal{U}'(\theta) = U_{c_0}(c_0(\theta), c_1(\theta))y_0'(\theta) + U_{c_1}(c_0(\theta), c_1(\theta))(Dk_0'(\theta) + y_1'(\theta)) \quad \forall \theta. \quad (24)$$

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

$$\max_{\{c_0(\theta), c_1(\theta), \mathcal{U}(\theta)\}} \int \omega(\theta) \mathcal{U}(\theta) dF(\theta) \quad (25)$$

s.t. $\mathcal{U}(\theta) = U(c_0(\theta), c_1(\theta))$, the resource constraint (19) and the incentive constraints (24).

A wealth tax. Alternatively, consider a tax $T_k(pk_1(\theta))$ on investors’ wealth in period 1.²³ 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

$$\mathcal{U}(\theta) \equiv U(p(k_0(\theta) - k_1(\theta)) + y_0(\theta), z_k(\theta) + y_1(\theta)) \geq U(p(k_0(\theta) - k_1(\hat{\theta})) + y_0(\theta), z_k(\hat{\theta}) + y_1(\theta))$$

for all $\theta, \hat{\theta}$. The local incentive constraints can therefore be written in the same general form as in the case of the asset sales tax, namely

$$\mathcal{U}'(\theta) = U_{c_0}(c_0(\theta), c_1(\theta))A(\theta) + U_{c_1}(c_0(\theta), c_1(\theta))B(\theta) \quad \forall \theta, \quad (26)$$

²³This 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.

with the only difference that, now, $A(\theta) = pk'_0(\theta) + y'_0(\theta)$ and $B(\theta) = y'_1(\theta)$. Hence, the second-best problem for the wealth tax is the same as above, except for the incentive constraints (26) instead of (24).²⁴

Other taxes. We can allow for other second-best tax instruments, such as consumption taxes, in an analogous way. Notably, we show in Supplementary Appendix C that the general incentive constraints (26) still apply, with modifications to the terms $A(\theta)$ and $B(\theta)$. 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 extends to any such second-best problem, including when combinations of taxes are available.²⁵

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). We use preferences $U(c_0, c_1) = G(C(c_0, c_1))$ where

$$C(c_0, c_1) = \left(c_0^{\frac{\sigma-1}{\sigma}} + \beta c_1^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{and} \quad G(C) = \frac{C^{1-\gamma}}{1-\gamma} \quad \text{with} \quad \sigma, \gamma > 0. \quad (27)$$

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.²⁶ To start with, we set $\sigma = 0.5$, $\gamma = 1$ (so $G(C) = \log(C)$) and $\beta = 0.5$.

As a baseline, we consider an asset price $\bar{p} = 1$ and dividends $\bar{D} = 2$ (so $\bar{D}/\bar{p} = 1/\beta$). 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 fixed at \bar{D} . Hence, this illustrates Special Case 1, an asset price change driven by a discount rate change.

Asset 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 redistribu-

²⁴In 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.

²⁵In Supplementary Appendix C, we characterize the Mirrlees (1971) Pareto optima for all these tax instruments and develop a numerical algorithm to compute them for changing asset prices.

²⁶(27) is a monotone transformation of the more standard intertemporally separable utility function $\sum_t \beta^t u(c_t)$ with $u(c) = c^{1-1/\sigma}/(1-1/\sigma)$. The reason for working with this monotone transformation is that, below, we will be interested in the limit as the intertemporal elasticity of substitution σ goes to zero, which is ill-defined for the standard specification. For example $c^{1-1/\sigma}/(1-1/\sigma) \rightarrow 0$ as $\sigma \rightarrow 0$ for all $c > 1$ (the numerator converges to zero and the denominator to $-\infty$). In contrast, (27) satisfies $U(c_0, c_1) \rightarrow G(\min\{c_0, c_1\})$ as $\sigma \rightarrow 0$, i.e., it converges to a (monotone transformation of a) Leontief utility function, as expected.




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.

tion runs from low- to high- θ types. As discussed above, asset sales x are naturally increasing in θ , so the optimum puts a tax on the (richer) buyers and a subsidy on the (poorer) sellers.

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.

Wealth tax. Figure 5 shows the respective 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 buy more assets and hence own more wealth at the beginning of the second period. Thus, in terms of levels, a progressive wealth tax with a positive tax burden on the rich and a subsidy on the poor 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.

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




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.

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.

In sum, 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 are 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 σ (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. Denote by $\Gamma^*(\sigma) \equiv \{(c_0^*(\theta, \sigma), c_1^*(\theta, \sigma))\}$ the optimal first-best allocation solving (21) subject to (19) when preferences are given by (27). We are interested in the limit as $\sigma \rightarrow 0$, so that $\mathcal{C}(c_0, c_1) = \min\{c_0, c_1\}$, which, as we show in Supplementary Appendix C, implies $c_0^*(\theta, 0) = c_1^*(\theta, 0) \equiv c^*(\theta)$. We need the following regularity assumption to obtain our convergence result:




FIGURE 6.—Capital gains tax with a decreasing intertemporal elasticity of substitution

ASSUMPTION 1: (i) *There exists a function α with $0 < \alpha(\theta) < 1$ for all θ such that*

$$c^{*'}(\theta) = \alpha(\theta)A(\theta) + (1 - \alpha(\theta))B(\theta)$$

and (ii) *letting $g(\theta) \equiv \log\left(\frac{1-\alpha(\theta)}{\beta\alpha(\theta)}\right)$, $g'(\theta)$ exists and is bounded for all θ .*

Part (i) amounts to a restriction on the Pareto weights $\omega(\theta)$ for any given heterogeneity $\{k_0(\theta), y_0(\theta), y_1(\theta)\}$. We show in Supplementary Appendix C that it follows from the first-best allocation $\Gamma^*(0)$ being incentive compatible when $\sigma = 0$, which is needed for the second-best allocation to be able to approach it. Part (ii) is a technical regularity condition on these Pareto weights.

Denote by $\Gamma^M(\sigma) \equiv \{(c_0^M(\theta, \sigma), c_1^M(\theta, \sigma))\}$ the solution to the Mirrlees problem (25) subject to (19) and (26) for the same Pareto weights $\{\omega(\theta)\}$. This yields the following result:

PROPOSITION 2: *Under Assumption 1 and utility (27), $\Gamma^M(\sigma) \xrightarrow[\sigma \rightarrow 0]{} \Gamma^*(\sigma)$ uniformly.*

Hence, for small σ and hence small distortions, our first-best results are informative about general second-best tax instruments, since the respective allocations are close.

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

Our analysis in this section thus far was based on a single-asset environment and emphasized the savings distortions from taxes. Therefore, it abstracted from taxes distorting portfolio choice. An important example of such portfolio distortions is the “lock-in” effect emphasized in the capital gains taxation literature (e.g. Holt and Shelton, 1962, Constantinides, 1983, Auerbach, 1991, Chari et al., 2005). Specifically, realization-based taxes incentivize deferring the liquidation of appreciated assets and thus distort optimal portfolio re-balancing in response to asset price changes.

We now show that an optimally designed second-best tax system does not introduce such distortions, even when it targets realized capital gains. The reason is that, in a multi-asset setting, it is always optimal to tax total *net* trades, i.e., to net all sales and purchases across the entire portfolio of assets, rather than taxing the gross gains from selling individual assets. As a result, when an investor sells one asset and uses the proceeds to purchase another one, there is generally no tax burden and therefore no lock-in effect.

As in Section 1, we introduce a second asset in the form of a bond. For simplicity, we consider a deterministic setup and therefore include a trading friction to prevent portfolio choice from being indeterminate. An investor's budget constraints are

$$c_0 = px - qb - \chi(x) + y_0 - T(x, b) \quad \text{and} \quad c_1 = D(k_0 - x) + b + y_1.$$

Trading frictions are captured by an adjustment cost $\chi(x)$ that is increasing in $|x|$ and convex. We allow for a general tax $T(x, b)$ on all trades x and b . In principle, such a tax could distort the investor's portfolio choice between capital and the bond, but we show that this is never optimal. To do so, we compare to a tax $T(z)$ on the total net trades $z \equiv px - qb - \chi(x)$, which by construction leaves the portfolio choice undistorted.²⁷

PROPOSITION 3: *Any optimum achieved by a tax $T(x, b)$ leaves portfolio choice undistorted and can be implemented with a tax on total net trades $T(z)$. Thus, there is no lock-in effect at the optimum.*

Hence, even if it is possible to condition taxes on individual (gross) trades, any optimum will not do so and instead will simply tax total net trades in or out of the portfolio. This result is akin to a production efficiency result (Diamond and Mirrlees, 1971). It implies that the second-best optimal policy does not introduce a lock-in effect. The lock-in effect results from the fact that, in practice, the capital gains from individual gross trades are taxed. Instead, by Proposition 3, pure portfolio re-balancing trades should not trigger a tax liability. This result is not specific to our setting, i.e. taxing net transactions eliminates the lock-in effect also in other settings.²⁸

²⁷Whether the tax is imposed in period 0 or 1 makes no difference for our argument. In particular, a tax on net trades in period 1, $b - Dx$, would be equivalent. We will relate our results to the deferral advantage, i.e., the interest advantage from deferring realization, in our general dynamic model in Section 5.

²⁸For example, Auerbach (1991) starts with a simple two-period illustration: "An investor, having accrued a first-period gain, g , must decide whether to realize the gain and reinvest at the rate of return, i , or hold the asset for an additional rate of return r . [A tax on realized capital gains...] makes the investor willing to hold even for a range of returns $r < i$." The investor's reinvestment decision is a case of pure portfolio re-balancing. Therefore not taxing such re-balancing eliminates the lock-in effect. Magnus (2024) makes a related proposal. Indeed, some real-world capital gains tax systems already include provisions that imply the taxation of net trades only. Examples are Section 1031 in the U.S. tax code on "like-kind exchanges" in the context of real estate investments, Section 351 on transfers to a corporation, and Section 368 on some types of mergers. Proposition 3 implies that such provisions are a good idea more broadly.

5. OPTIMAL TAXATION IN THE GENERAL MODEL

We now return to the case of first-best tax instruments and show how our results on optimal taxation from the deterministic two-period model extend to the general model in Section 1.

5.1. Risk and Borrowing

First Best. With lump-sum taxes, the investors’ sequential budget constraints (2) become

$$\begin{aligned} c_t(s^t, \theta) + p_t(s^t)(k_{t+1}(s^t, \theta) - k_t(s^{t-1}, \theta)) + q_t(s^t)b_{t+1}(s^t, \theta) \\ = y_t(\theta) + D_t(s^t)k_t(s^{t-1}, \theta) + b_t(s^{t-1}, \theta) - T_t(\theta, s^t) \quad \forall t, s^t. \end{aligned}$$

We allow taxes and transfers $T_t(\theta, s^t)$ to be indexed by s^t . In order to ensure that risk is relevant, we assume

$$\int T_t(\theta, s^t)dF(\theta) = 0 \quad \forall t, s^t,$$

so the economy cannot insure itself other than through trading capital and the bond with the rest of the world. In other words, it does not have access to the full set of Arrow-Debreu insurance markets. The first-best allocation is the solution to

$$\max_{\{c_t(\theta, s^t), C_t(s^t), K_{t+1}(s^t), B_{t+1}(s^t)\}} \int \omega(\theta)U(\{c_t(\theta, s^t)\})dF(\theta) \text{ s.t. } \int c_t(\theta, s^t)dF(\theta) = C_t(s^t) \quad \forall t, s^t$$

and the aggregate resource constraint (3).²⁹

Shocks to asset prices and dividends. We are now in position to revisit how shocks to asset prices and cash flows induce changes in the optimal tax burden. We begin with a pairwise comparison across two arbitrary potential histories, mirroring the preceding comparative static exercise. Let s^t and \bar{s}^t denote two alternative histories starting from a common s_0 . We refer to \bar{s}^t as the “baseline” and s^t as the alternative history of interest. These histories could feature different stochastic discount factors ($m_t(s^t) \neq m_t(\bar{s}^t)$) or dividend streams ($D_t(s^t) \neq D_t(\bar{s}^t)$), and hence represent different prices p_t and q_t .

To fix ideas, consider Figure 7: the baseline is a steady state with constant dividend $D_t(\bar{s}^t) = \bar{D}$, discount factor $m_t(\bar{s}^t) = 1/\bar{R}$, and asset prices $p_t(\bar{s}^t) = \bar{p}$ and $q_t(\bar{s}^t) = \bar{q}$.³⁰ At time $t = 1$, an alternative path is realized and $D_t(s^t), m_t(s^t), p_t(s^t), q_t(s^t)$ deviate from the initial steady state. The question we consider is: how does the optimal tax system redistribute in response to this different realization? For example, as in Figure 7, $p_t(s^t)$ starts increasing but without

²⁹Without loss of generality, we assume that the government does not own bonds or capital directly, and thus the aggregates B_t and K_t reflect the aggregated holdings of private investors.

³⁰An alternative baseline is a BGP where dividends grow at some constant rate G , so $D_t(\bar{s}^t) = G^t \bar{D}_0$, and rates of return are constant $m_t(\bar{s}^t) = 1/\bar{R}$. This could be the BGP of an equilibrium model with a neoclassical production side and productivity growth of the type discussed in Section 1.5.




FIGURE 7.—Example time paths for asset price, dividends, and discount rate

any corresponding increase in cash flows $D_t(s^t)$ (panels (a) and (b)); equivalently, the asset discount rate jumps up initially and then declines secularly to a lower long-run level (panel (c)), which generates capital gains. The example in Figure 7 therefore corresponds to Special Case 1 from Section 1.3 with an asset price change driven entirely by a discount rate change.

Optimal taxation. To simplify notation, we suppress s^t and simply write $\Delta p_t = p_t(s^t) - p_t(\bar{s}^t)$ and analogously for ΔD_t , Δq_t and $\Delta T_t(\theta)$. Similarly, we write $k_t(\theta)$ as a shorthand for $k_t(\theta, s^{t-1})$, $b_t(\theta)$ for $b_t(\theta, s^{t-1})$, $x_t(\theta)$ for $x_t(\theta, s^t)$, and analogously for the corresponding aggregates K_t , B_t , and X_t . The next result extends Proposition 1 to the general model:

PROPOSITION 4: *Consider a shock that changes asset prices by $\{\Delta p_t, \Delta q_t\}$ and dividends by $\{\Delta D_t\}$. Then the following tax change is an optimal response for all s^t, \bar{s}^t :*

$$\Delta T_t(\theta) = x_t(\theta)\Delta p_t + k_t(\theta)\Delta D_t - b_{t+1}(\theta)\Delta q_t - \Omega(\theta)(X_t\Delta p_t + K_t\Delta D_t - B_{t+1}\Delta q_t).$$

Compared to Proposition 1, the only difference is the additional compensation for changes in the bond price Δq_t . 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. In particular, when the interest rate increases ($\Delta q_t < 0$), Proposition 4 prescribes higher taxes on savers ($b_{t+1}(\theta) > 0$) and lower taxes on borrowers ($b_{t+1}(\theta) < 0$), which could be achieved by a standard income tax (with interest deductibility in case of debt). As described in Section 2, this could be driven by heterogeneity in the underlying income profiles, with borrowers, for example, featuring a more backloaded income process (e.g., due to more future human capital).

More generally, a noteworthy example to illustrate Proposition 4 is a shock to the pricing kernel Δm_t such that $\mathbb{E}_t[\Delta m_{t+1}] = 0$. Then, by (4), $\Delta q_t = 0$, so the risk-free rate is unchanged. If dividends are also held fixed, this corresponds to a pure risk-premium change. In this case,

Proposition 4 collapses to

$$\Delta T_t(\theta) = x_t(\theta)\Delta p_t - \Omega(\theta)X_t\Delta p_t \quad (28)$$

just like in 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.

As discussed in Section 1.3, such a risk premium change could equivalently be driven by changes in subject beliefs about future cash flows because, by equations (4) and (5), the asset prices $\{p_t, q_t\}$ depend on the probabilities $\pi(s^t)$ with which different histories occur (through the expectations operator). Thus, if the shock increases subjective probabilities $\tilde{\pi}(s^t)$ on histories in which cash flows $D_t(s^t)$ are high, without affecting actual dividends and while leaving $\mathbb{E}_t[m_{t+1}]$ unchanged (so that $\Delta q_t = 0$), the optimal tax response also collapses back to (28).

As in our two-period model, the timing of taxes is not pinned down; Proposition 4 picks one particular normalization, which taxes or subsidizes the trading gains and losses as well as the change in dividend and interest income relative to the baseline period by period. Our next result derives a present-value condition that any tax implementation needs to satisfy. Our condition concerns the difference in the present-discounted value of taxes in any fixed history s^t compared to the present value over *all* possible baseline histories \tilde{s}^t :

$$\begin{aligned} \mathbb{E}_0 \sum_{t=0}^T m_{0 \rightarrow t}(\tilde{s}^t) \Delta T_t(\theta) &= \mathbb{E}_0 \sum_{t=0}^T m_{0 \rightarrow t}(\tilde{s}^t) [T_t(\theta, s^t) - T_t(\theta, \tilde{s}^t)] \\ &= \sum_{t=0}^T q_{0 \rightarrow t} T_t(\theta, s^t) - \mathbb{E}_0 \sum_{t=0}^T m_{0 \rightarrow t}(\tilde{s}^t) T_t(\theta, \tilde{s}^t), \end{aligned}$$

where the expectation \mathbb{E}_0 is taken over all possible reference histories \tilde{s}^t and the second line uses the pricing condition for the risk-free rate.³¹ That is, we consider the change in the present value of taxes along a particular history s^t relative to the unconditional present value of taxes.

PROPOSITION 5: *Consider a history s^t that differs relative to any alternative history \tilde{s}^t in asset prices by $\{\Delta p_t, \Delta q_t\}$ and dividends by $\{\Delta D_t\}$, where $\Delta p_t = p_t(s^t) - p_t(\tilde{s}^t)$ and similarly for q_t and D_t . Then, letting \mathbb{E}_0 represent expectation over \tilde{s}^t , the relative tax burden of s^t is such that*

$$\mathbb{E}_0 \left[\sum_{t=0}^T m_{0 \rightarrow t}(\tilde{s}^t) \Delta T_t(\theta) \right]$$

³¹ $q_{0 \rightarrow t} \equiv \mathbb{E}_0[m_{0 \rightarrow t}(\tilde{s}^t)]$ is the value of a risk-free bond at time $t = 0$ that pays off one unit of consumption at time t in all possible states \tilde{s}^t .

$$= \mathbb{E}_0 \left[\sum_{t=0}^T m_{0 \rightarrow t}(\tilde{s}^t) [x_t(\theta) \Delta p_t + k_t(\theta) \Delta D_t - b_{t+1}(\theta) \Delta q_t - \Omega(\theta) (X_t \Delta p_t + K_t \Delta D_t - B_{t+1} \Delta q_t)] \right].$$

Thus, our results from the two-period case generalize to the expected present value of taxes.

Role of the government. Proposition 5 speaks to the fact that there are many tax schemes that implement the same optimal allocation, with private agent portfolios adjusting accordingly to maintain the same consumption allocation. In fact, it is possible to set $\Delta T_t(\theta) = 0$ for all $t \geq 1$ by ensuring, through appropriate taxation in period 0, that all investors θ hold the market portfolio from period 1 on, i.e.

$$k_t(\theta, s^{t-1}) = \Omega(\theta) K_t(s^{t-1}) \quad \text{and} \quad b_t(\theta, s^{t-1}) = \Omega(\theta) B_t(s^{t-1}) \quad \forall t \geq 1, s^t.$$

In this case, there is no trade within the economy, only with the rest of the world, so all investors are equally affected by future shocks. As a result, there is no scope for redistributive taxation going forward, and $T_t(\theta, s^t) = 0$ for $t > 0$ along *any* history s^t .³² In Supplementary Appendix D, we use a two-period variant of our general model with risk to explain this in more detail.

Why do we instead focus on tax implementations beyond the initial period? The first reason is that we regard our preferred implementation, which targets realized capital gains and dividends in each period, as much simpler. Tracking cash flows and trades along the way is easier than assessing the stocks of all portfolio holdings in the initial period. Second, in practice, it may not be feasible to ensure that all investors hold the full market portfolio after the initial period. For example, and outside our model, some assets (e.g., a startup) may not be publicly traded, rendering an implementation that redistributes only in period zero infeasible. However, when the founder takes the company public and sells part of their shares, our preferred implementation will tax (or subsidize) those sales when they happen, which is straightforward.

Alternative baseline asset prices. As discussed above, our tax formulas are based on a comparison between alternative price and dividend paths, compensating investors for a given shock relative to a baseline path. As a result, the asset price change Δp_t refers to the (contemporaneous) difference between $p_t(s^t)$ relative to the baseline price $p_t(\bar{s}^t)$, not a price increment $p_t - p_{t-1}$. Notably, the baseline price $p_t(\bar{s}^t)$ does not necessarily correspond to the historical acquisition prices at which the investor purchased the asset. In other words, the realized capital gains $x_t(\theta) \Delta p_t$ in (28) may differ from taxable capital gains in real-world capital gains tax systems, which are typically based on price changes between acquisition and sale.

An interesting alternative problem would be to characterize optimal taxes subject to the institutional constraint that capital gains must be computed relative to the historical purchase price.

³²To focus on asset prices and capital gains, we have assumed that the income paths $\{y_t(\theta)\}_{t=0}^T$ are deterministic. With stochastic income paths $y_t(\theta, s^t)_{t=0}^T$, taxes would still be required beyond the initial period to insure income risk, taking the form $T_t(\theta, s^t) = y_t(\theta, s^t) - \Omega(\theta) Y_t(s^t)$.

In general, this would not fully achieve the compensation we focus on here (and only apply to asset sales, not purchases), but it may be able to approximate it, giving rise to a different, Ramsey-style problem. It would also introduce technical complications because the optimal policy would depend on the history of past prices, rather than just contemporaneous comparisons (see, e.g., [Boerma et al. \(2023\)](#), for tools to solve such problems). We leave this for future research.

5.2. *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 by borrowing against them. [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. [...] 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.”

Our analysis in the preceding subsection is useful for determining how optimal taxation should treat borrowing versus selling. The most instructive case is that of positive capital gains $\Delta p_t > 0$ but without a corresponding change in interest rates $\Delta q_t = 0$, which is the case of a pure risk-premium change—see equation (28). Perhaps surprisingly, the tax formula is *independent of* whether and how much investors borrow when their assets appreciate and is, in fact, identical to Corollary 1. Contrary to a recent proposal by [Fox and Liscow \(2024\)](#), it is not necessary to tax borrowing.

The intuition (sometimes missed in the popular debate) is that, also with the option to borrow, investors need to sell their appreciating assets *at some point* in order to repay their loans and benefit from rising asset prices. 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.³³ On the other hand, if investors do sell to repay the loan, the realized trade should be taxed at that point.

[The Economist \(2024\)](#) quote above 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 model here does not allow for this possibility, Section 6.2 considers a setup with

³³An exception is “stepped-up basis” which we discuss in Section 6.3. Moreover, investors may benefit from capital gains even without selling if a collateral constraint gets relaxed by a rising asset price ([Fagereng et al., forthcoming](#)). Here, given our focus on the top of the wealth distribution, we abstract from such binding constraints.

heterogeneous returns with precisely this feature – see equation (32) below. Still, it turns out that the optimal tax formula is unaffected. The intuition is that, while such return differences are undoubtedly important, they are not specific 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 an orthogonal issue that should not be considered tax avoidance.

5.3. The deterministic case

The case without uncertainty is again particularly instructive because the two assets collapse to a single one and we can think of Proposition 5 as a comparative static exercise comparing taxes under two different time paths for discount rates and dividends $\{R_t, D_t\}_{t=0}^T$ and associated asset prices $\{p_t\}_{t=0}^T$ satisfying (7). Since $\bar{m}_{0 \rightarrow t} = \bar{R}_{0 \rightarrow t}^{-1}$, we obtain the following corollary of Proposition 5:

COROLLARY 3: *Suppose asset prices change by $\{\Delta p_t\}$ and dividends by $\{\Delta D_t\}$. Then optimal taxes $\{T_t(\theta)\}$ change such that*

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

Proposition 5 now applies to the comparison between any two price and dividend paths.³⁴ In Special Case 1, where asset prices change exclusively because of a change in discount rates, i.e. $\Delta D_t = 0$ for all t , we see that optimal redistributive taxes condition only on realized trades $\{x_t(\theta), X_t\}$ and not on asset holdings $\{k_t(\theta), K_t\}$.³⁵

In Special Case 2, where asset prices change exclusively because of a change in future dividends, the returns \bar{R}_{t+1} remain unchanged for all $t \geq 1$ and the asset price change satisfies (9). 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 that are instead driven exclusively by a change in future dividends $\{\Delta D_t\}$. Then we obtain the following result:

³⁴Corollary 3 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

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

³⁵Note also that Corollary 3 fixes the present value of taxes, consistent with a requirement of Vickrey (1939) for desirable tax systems: “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.”

COROLLARY 4: *Suppose the change in prices $\{\Delta p_t\}$ is exclusively driven by the change in dividends $\{\Delta D_t\}$. Then optimal taxes change such that*

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

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 $\Delta T_0(\theta) = [k_0(\theta) - \Omega(\theta)K_0] (\Delta D_0 + \Delta p_0)$ and $\Delta T_t = 0, t \geq 1$, i.e. a one-time accrual-based capital gains tax at $t = 0$.

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 (9) 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

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

Thus, a period-by-period implementation would set $\Delta T_t(\theta) = k_0(\theta)\Delta D_t - \Omega(\theta)K_0\Delta D_t$ for all t , which does *not* correspond to a tax on accrued gains (nor dividend income) in each period.

6. EXTENSIONS

In this section, we show how the results derived in the benchmark setting extend to richer environments, namely a closed economy general equilibrium model, heterogeneous returns, 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

$$\int k_0(\theta) dF(\theta) = \int k_1(\theta) dF(\theta) = K. \quad (29)$$

The asset price p must adjust to satisfy the market clearing condition (29). For example, if preferences are given by (27), the equilibrium asset price p^* can be solved in closed form:

$$p^* = \beta D \left(\frac{Y_0}{Y_1 + DK} \right)^{\frac{1}{\sigma}}. \quad (30)$$

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 increases the asset price p^* proportionally (holding dividends fixed). More generally, regardless of what causes the change in the equilibrium price p^* , we obtain the following result:

PROPOSITION 6: *Suppose the equilibrium asset price changes from \bar{p}^* to $p^* = \bar{p}^* + \Delta p^*$, holding dividends D and the aggregate endowment (Y_0, Y_1) 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 dividends and the aggregate endowment are held fixed. Hence, the 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 via Slutsky compensation as in Lemma 2.

We can also consider a change in dividends D in general equilibrium. Equation (30) 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 lent 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.

Since our first-best exercise here considers lump-sum taxes that do not distort investor behavior, the equilibrium asset price in equation (30) is independent of the level of taxes $T_0(\theta), T_1(\theta)$. This result would not survive in the presence of distortive taxes, as in Section 4, because equilibrium asset prices may then depend on taxes and transfers. More generally, in a general-equilibrium version of our dynamic model with risk from Section 5, distortive taxation would generally affect both the stochastic discount factor and dividends. Thus, optimal taxes and asset prices would need to be jointly determined as the solution to a fixed point problem; see, e.g., the general-equilibrium optimality formulas developed by Rothschild and Scheuer (2013, 2014, 2016) and Schulz et al. (2022) in the context of labor income taxation. Embedding our second-best analysis from Section 4 in general equilibrium would be a promising first step in this direction.

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 $R(\theta) \equiv D(\theta)/p$.³⁶

³⁶See, for example, Gerritsen et al. (2020), Schulz (2021) and Guvenen et al. (2023, 2024).

Just introducing this additional heterogeneity into our baseline model does not change our results on Pareto optimal tax policy. The reason is that, in the absence of further frictions, it is efficient 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 $R = D^{\max}/p$, effectively returning us 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, we build on our analysis in Section 4.4 and re-introduce a bond and some trading friction. Using the same notation as there, 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 adjustment cost $\chi(x)$ ensures that it is no longer efficient to allocate all capital to the individual with the highest return. An investor’s optimal asset sales $x(\theta)$ satisfies

$$qD(\theta) + \chi'(x(\theta)) = p. \quad (31)$$

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. Due to the convex adjustment cost, investors with higher returns $D(\theta)$ will sell less and hold more of the asset. Also note that the presence of adjustment costs in (31) implies that

$$R(\theta) \equiv \frac{D(\theta)}{p} \geq \frac{1}{q}, \quad (32)$$

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 different investors’ returns $R(\theta)$ changing differentially in response to heterogeneous cash flow changes.

The aggregate resource constraint can be written as

$$\int c_0(\theta)dF(\theta) + q \int c_1(\theta)dF(\theta) = Y \quad (33)$$

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 problem takes the same form as in Section 3. Normalizing $T_1(\theta) = 0$:

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*

$$T_0(\theta) = \bar{T}_0(\theta) + x(\theta)\Delta p - \Omega(\theta)X\Delta p + \mathcal{O}(\Delta x(\theta)^2) \quad \text{where } \Delta x(\theta) \equiv x(\theta) - \bar{x}(\theta).$$

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 (31) 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. As already anticipated above, in equilibrium, different investors experience differential returns given by

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

Consider now an increase in dividends $D(\theta)$ for *some subset* of the investors in the economy. This induces an increase in the equilibrium asset price p^* for *all* investors, including for those whose dividends did not change.³⁷ 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 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 satisfies $\Delta T_0(\theta) \approx x(\theta)\Delta p^*$ as in Propositions 6 and 7.

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

³⁷If bonds are in fixed supply as well, then the bond price q also adjusts, but in general this does not undo the change in the average dividends.

in the baseline model where everyone achieves the same rate of return. With heterogeneous dividends, these effects no longer cancel (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. *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.³⁸

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

$$V_t(\theta) = U(c_t(\theta), \dots, c_{t+\tau-1}(\theta)) + \alpha\beta^\tau V_{t+\tau}(\theta), \quad (34)$$

where $0 \leq \alpha \leq 1$ measures altruism toward the next generation and $U(c_t, \dots, c_{t+\tau-1})$ is homothetic. The sequential budget constraint is still given by (10) 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.³⁹ The Pareto problem is to maximize $\int \omega(\theta)V_0(\theta)dF(\theta)$ subject to (3) where $\omega(\theta)$ is the Pareto weight on dynasty θ . Hence, everything collapses to the multi-period model from Section 5 and Corollary 3 applies with the only modification that, in general, it includes 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) Corollary 3 still applies, so does the discussion in Section 3.3 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 at death

³⁸Many 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.”

³⁹Repeated substitution of (34) implies that the dynasty θ ’s utility at time 0 is given by

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

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. From the point of a view of a dynasty or the social planner, 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. A tax avoidance strategy of wealthy families known as “buy, borrow, die” has received attention in recent years (e.g. Ensign and Rubin, 2021, *The Economist*, 2024).⁴⁰ The idea is to borrow against appreciating assets rather than selling them and then taking advantage of the stepped-up basis at death, thereby avoiding capital gains taxes altogether. In combination with Section 5.2, 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 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.

It is useful to juxtapose our results 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. We show that this intuition is, in general, incorrect. Instead, we show that optimal taxes can always be implemented in such a way that they 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 taxing asset holdings each period is generally suboptimal.

REFERENCES

- ADAM, KLAUS, ALBERT MARCET, AND JOHANNES BEUTEL (2017): “Stock Price Booms and Expected Capital Gains,” *American Economic Review*, 107 (8), 2352–2408. [9]
- AGERSNAP, OLE AND OWEN ZIDAR (2021): “The Tax Elasticity of Capital Gains and Revenue-Maximizing Rates,” *American Economic Review: Insights*, 3 (4), 399–416. [6]
- ATKESON, ANDREW, JONATHAN HEATHCOTE, AND FABRIZIO PERRI (2024): “There is No Excess Volatility Puzzle,” *NBER Working Paper 32481*. [8]
- ATKINSON, ANOTHONY AND JOSEPH STIGLITZ (1976): “The design of tax structure: Direct versus indirect taxation,” *Journal of Public Economics*, 6, 55–75. [1, 11]
- AUERBACH, ALAN (1989): “Capital Gains Taxation and Tax Reform,” *National Tax Journal*, 42 (3), 391–401. [1, 2]

⁴⁰Fox and Liscow (2025) question this strategy’s prevalence and argue that “buy, save, die” is more widespread.

- (1991): "Retrospective Capital Gains Taxation," *American Economic Review*, 81 (1), 167–178. [29, 30]
- AZEVEDO, EDUARDO M. FLORIAN SCHEUER, KENT SMETTERS, AND MIN YANG (2025): "Dilution vs. Risk Taking: Capital Gains Taxes and Entrepreneurship," NBER Working Paper 34512, National Bureau of Economic Research. [4]
- BANKS, JAMES AND PETER DIAMOND (2010): "The base for direct taxation," in *Dimensions of Tax Design: The Mirrlees Review*, ed. by Stuart Adam, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba, Oxford University Press, 548–648. [5]
- BARRO, ROBERT AND GARY BECKER (1989): "Fertility Choice in a Model of Economic Growth," *Econometrica*, 57 (2), 481–501. [41]
- BASTANI, SPENCER AND DANIEL WALDENSTROM (2020): "How Should Capital Be Taxed?" *Journal of Economic Surveys*, 34 (4), 812–846. [5]
- BENHABIB, JESS AND BÁLINT SZÖKE (2021): "Optimal Positive Capital Taxes at Interior Steady States," *American Economic Journal: Macroeconomics*, 13 (1), 114–150. [12]
- BOERMA, JOB, GEORGI RIABOV, AND ALEH TSYVINSKI (2023): "Policy with Stochastic Hysteresis," *Mimeo*. [35]
- BONNET, ODRAN, PIERRE-HENRI BONO, GUILLAUME FLAMERIE DE LA CHAPELLE, AND ETIENNE WASMER (2014): "Does housing capital contribute to inequality? A comment on Thomas Piketty's Capital in the 21st Century," *SciencePo Working paper 2014-07*. [6]
- BORDALO, PEDRO, NICOLA GENNAIOLI, RAFAEL LA PORTA, MATTHEW O'BRIEN, AND ANDREI SHLEIFER (2023): "Long-Term Expectations and Aggregate Fluctuations," *NBER Macroeconomics Annual*, 38. [9]
- CAMPBELL, JOHN (2018): *Financial Decisions and Markets*, Princeton University Press. [8]
- CAMPBELL, JOHN AND ROBERT SHILLER (1988): "The dividend-price ratio and expectations of future dividends and discount factors," *The Review of Financial Studies*, 1 (3), 195–228. [1, 8]
- CARAMP, NICOLAS AND DEJANIR SILVA (2023): "Monetary Policy and Wealth Effects: The Role of Risk and Heterogeneity," *UC Davis Working Paper*. [17]
- CATHERINE, SYLVAIN, MAX MILLER, JAMES PARON, AND NATASHA SARIN (2024): "Interest-rate Risk and Household Portfolios," Tech. rep., Working Paper. [6]
- CATHERINE, SYLVAIN, MAX MILLER, AND NATASHA SARIN (2020): "Social security and trends in wealth inequality," *Working Paper*. [6]
- CHAMLEY, CHRISTOPHE (1986): "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54 (3), 607–622. [1, 11]
- CHARI, V.V. AND PATRICK J. KEHOE (1999): "Chapter 26 Optimal fiscal and monetary policy," in *Handbook of Macroeconomics*, Elsevier, vol. 1 of *Handbook of Macroeconomics*, 1671–1745. [12]
- CHARI, V.V., JUAN PABLO NICOLINI, AND PEDRO TELES (2018): "Optimal Capital Taxation Revisited," Working Papers 752, Federal Reserve Bank of Minneapolis. [12]
- (2020): "Optimal capital taxation revisited," *Journal of Monetary Economics*, 116 (C), 147–165. [12]
- CHARI, V. V., MIKHAIL GOLOSOV, AND ALEH TSYVINSKI (2005): "Business Start-ups, The Lock-in Effect, and Capital Gains Taxation," Tech. rep., Working Paper. [6, 29]
- CIOFFI, RICCARDO (2021): "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality," *Princeton University Working Paper*. [6]
- COCHRANE, JOHN (2011): "Presidential Address: Discount Rates," *Journal of Finance*, 66 (4), 1047–1108. [1]
- CONSTANTINIDES, GEORGE (1983): "Capital Market Equilibrium with Personal Tax," *Econometrica*, 51 (3), 611–636. [6, 29]
- COVEN, JOSHUA, SEBASTIAN GOLDER, ARPIT GUPTA, AND ABDOULAYE NDIAYE (2024): "Property Taxes and Housing Allocation Under Financial Constraints," CEPR Discussion Papers 19230, C.E.P.R. Discussion Papers. [6]
- DÁVILA, EDUARDO AND ANTON KORINEK (2018): "Pecuniary externalities in economies with financial frictions," *The Review of Economic Studies*, 85 (1), 352–395. [15]
- DIAMOND, PETER AND JAMES MIRRLEES (1971): "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review*, 61, 8–27. [30]
- ENSIGN, RACHEL LOUISE AND RICHARD RUBIN (2021): "Buy, Borrow, Die: How Rich Americans Live Off Their Paper Wealth," *Wall Street Journal*. [42]
- EPSTEIN, LARRY AND STANLEY ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57 (4), 937. [7]
- FAGERENG, ANDREAS, MATTHIEU GOMEZ, EMILIE GOUIN-BONENFANT, MARTIN BLOMHOFF HOLM, BENJAMIN MOLL, AND GISLE NATVIK (forthcoming): "Asset-Price Redistribution," *Journal of Political Economy*. [6, 15, 35]
- FARHI, EMMANUEL, CHRISTOPHER SLEET, IVÁN WERNING, AND SEVIN YELTEKIN (2012): "Non-linear Capital Taxation without Commitment," *Review of Economic Studies*, 79, 1469–1493. [12]

- FELDSTEIN, MARTIN, JOEL SLEMROD, AND SHLOMO YITZHAKI (1980): “The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains,” *The Quarterly Journal of Economics*, 94 (4), 777–791. [6]
- FEREY, ANTOINE, BEN LOCKWOOD, AND DMITRY TAUBINSKY (2024): “Sufficient Statistics for Nonlinear Tax Systems with General Across-Income Heterogeneity,” *American Economic Review*, forthcoming. [11]
- FOX, EDWARD AND ZACHARY LISCOW (2024): “No More Tax-Free Lunch for Billionaires: Closing the Borrowing Loophole,” *Tax Notes Federal*. [1, 35]
- (2025): “The Role of Unrealized Gains and Borrowing in the Taxation of the Rich,” *Mimeo*. [42]
- GERRITSEN, AART, BAS JACOBS, ALEXANDRA VICTORIA RUSU, AND KEVIN SPIRITUS (2020): “Optimal Taxation of Capital Income with Heterogeneous Rates of Return,” *CESifo working paper 8395*. [38]
- GOLOSOV, MIKHAIL, ALEH TSYVINSKI, AND IVÁN WERNING (2007): “New Dynamic Public Finance: A User’s Guide,” *NBER Macroeconomics Annual*, 21, 317–388. [5, 11]
- GOMEZ, MATTHIEU (2016): “Asset prices and wealth inequality,” *Working paper*. [6]
- GOMEZ, MATTHIEU AND EMILIEN GOUIN-BONENFANT (2020): “A q-theory of inequality,” *Working paper*. [6]
- GORDON, MYRON AND ELI SHAPIRO (1956): “Capital equipment analysis: the required rate of profit,” *Management science*, 3 (1), 102–110. [9]
- GREENWALD, DANIEL, MATTEO LEOMBRONI, HANNO LUSTIG, AND STIJN VAN NIEUWERBURGH (2021): “Financial and total wealth inequality with declining interest rates,” *NBER working paper 28613*. [6]
- GREENWALD, DANIEL, MARTIN LETTAU, AND SYDNEY LUDVIGSON (2019): “How the wealth was won: Factors shares as market fundamentals,” *NBER working paper 25769*. [1]
- GUVENEN, FATIH, GUEORGUI KAMBOUROV, BURHAN KURUSCU, AND SERGIO OCAMPO (2024): “Book-Value Wealth Taxation, Capital Income Taxation, and Innovation,” *NBER working paper 32585*. [12, 38]
- GUVENEN, FATIH, GUEORGUI KAMBOUROV, BURHAN KURUSCU, SERGIO OCAMPO, AND DAPHNE CHEN (2023): “Use It or Lose It: Efficiency and Redistributive Effects of Wealth Taxation,” *The Quarterly Journal of Economics*, 138 (2), 835–894. [12, 38]
- HAIG, ROBERT (1921): “The Concept of Income – Economic and Legal Aspects,” in *The Federal Income Tax*, ed. by Robert M. Haig, New York: Columbia University Press, 1–28. [1, 6, 11, 21]
- HELLWIG, CHRISTIAN AND NICOLAS WERQUIN (2024): “Using Consumption Data to Derive Optimal Income and Capital Tax Rates,” *Working Paper*. [11]
- HOLT, CHARLES C. AND JOHN P. SHELTON (1962): “The Lock-in Effect of the Capital Gains Tax,” *National Tax Journal*, 15 (4), 337–352. [29]
- JERMANN, URBAN (1998): “Asset pricing in production economies,” *Journal of Monetary Economics*, 41 (2), 257–275. [12]
- JUDD, KENNETH (1985): “Redistributive taxation in a simple perfect foresight model,” *Journal of Public Economics*, 28 (1), 59–83. [1, 12]
- KALDOR, NICHOLAS (1955): *An Expenditure Tax*, London and New York, relevant excerpt available at <https://benjaminmoll.com/kaldor/>: Routledge. [6, 24]
- KUHN, MORITZ, MORITZ SCHULARICK, AND ULRIKE STEINS (2020): “Income and wealth inequality in America, 1949–2016,” *Journal of Political Economy*, 128 (9), 3469–3519. [6]
- LARRAIN, BORJA AND MOTOHIRO YOGO (2008): “Does firm value move too much to be justified by subsequent changes in cash flow,” *Journal of Financial Economics*, 87 (1), 200–226. [8]
- LEISERSON, GREG AND DANNY YAGAN (2021): “What Is the Average Federal Individual Income Tax Rate on the Wealthiest Americans?” *U.S. Council of Economic Advisors, Washington, DC, White House*. [1]
- LUCAS, ROBERT (1990): “Supply-Side Economics: An Analytical Review,” *Oxford Economic Papers*, 42 (2), 293–316. [1]
- MAGNUS, GIDEON (2024): “Sensible Taxation of Investment Returns,” *Tax Notes Federal*. [30]
- MARTÍNEZ-TOLEDANO, CLARA (2022): “House Price Cycles, Wealth Inequality and Portfolio Reshuffling,” *Working Paper, Imperial College London*. [6]
- MAS-COLELL, ANDREU, MICHAEL WHINSTON, AND JERRY GREEN (1995): *Microeconomic Theory*, New York: Oxford University Press. [17]
- MIRRELES, JAMES (1971): “An Exploration in the Theory of Optimum Income Taxation,” *The Review of Economic Studies*, 38 (2), 175–208. [2, 4, 6, 26]
- MOLL, BENJAMIN (2020): “Comment on “Sources of U.S. Wealth Inequality: Past, Present, and Future”,” *NBER Macroeconomics Annual*, 35. [6, 15]
- MSALL, LUCY AND OLE-ANDREAS NÆSS (2025): “Never-Realized Capital Gains,” Tech. rep., University of Chicago. [6]
- NAGEL, STEFAN (2024): “Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions,” Tech. rep., University of Chicago. [8]
- OECD (2021): “Inheritance Taxation in OECD Countries,” *OECD Tax Policy Studies*. [42]

- PAISH, FRANK (1940): "Capital Value and Income," *Economica*, 7 (28), 416–418. [6]
- PIKETTY, THOMAS, EMMANUEL SAEZ, AND GABRIEL ZUCMAN (2023): "Rethinking capital and wealth taxation," *Oxford Review of Economic Policy*, 39 (3), 575–591. [1]
- POTERBA, JAMES (2002): "Taxation, risk-taking, and household portfolio behavior," in *Handbook of Public Economics*, ed. by A. J. Auerbach and M. Feldstein, Elsevier, vol. 3 of *Handbook of Public Economics*, chap. 17, 1109–1171. [6]
- ROGNLIE, MATTHEW (2015): "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?" *Brookings Papers on Economic Activity*, 46 (1 (Spring)), 1–69. [6]
- ROTHSCHILD, CASEY AND FLORIAN SCHEUER (2013): "Redistributive Taxation in the Roy Model," *Quarterly Journal of Economics*, 128, 623–668. [38]
- (2014): "A Theory of Income Taxation under Multidimensional Skill Heterogeneity," *NBER Working Paper No. 19822*. [38]
- (2016): "Optimal Taxation with Rent-Seeking," *Review of Economic Studies*, 83, 1225–1262. [38]
- SAEZ, EMMANUEL (2002): "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes," *Journal of Public Economics*, 83, 217–230. [11]
- SAEZ, EMMANUEL AND STEFANIE STANTCHEVA (2018): "A simpler theory of optimal capital taxation," *Journal of Public Economics*, 162 (C), 120–142. [11]
- SAEZ, EMMANUEL, DANNY YAGAN, AND GABRIEL ZUCMAN (2021): "Capital Gains Withholding," *UC Berkeley working paper*. [1]
- SCHEUER, FLORIAN AND JOEL SLEMROD (2021): "Taxing Our Wealth," *Journal of Economic Perspectives*, 35 (1), 207–230. [5]
- SCHEUER, FLORIAN AND ALEXANDER WOLITZKY (2016): "Capital Taxation under Political Constraints," *American Economic Review*, 106, 2304–2328. [11]
- SCHULZ, KARL (2021): "Redistribution of Return Inequality," *Working paper*. [38]
- SCHULZ, KARL, ALEH TSYVINSKI, AND NICOLAS WERQUIN (2022): "Generalized Compensation Principle," *Theoretical Economics*, 18, 1665–1710. [38]
- SHILLER, ROBERT (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71 (3), 421–436. [1]
- SHOURIDEH, ALI (2012): "Taxation of Wealthy Individuals," *Working Paper*. [12]
- SIMONS, HENRY (1938): *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy*, University of Chicago Press Chicago, Ill. [1, 6, 11, 21]
- SMITH, KATE AND HELEN MILLER (2023): "It's all about the base: Taxing business owner-managers," *Working Paper*. [6]
- STANTCHEVA, STEFANIE (2020): "Dynamic Taxation," *Annual Review of Economics*, 12, 801–831. [5]
- STRAUB, LUDWIG AND IVÁN WERNING (2020): "Positive Long-Run Capital Taxation: Chamley-Judd Revisited," *American Economic Review*, 110 (1), 86–119. [12]
- THE ECONOMIST (2024): "Sweet gains, bro," *June 22 2024 Issue, Finance and Economics Section*. [35, 42]
- TRUST AND WILL (2024): "How the Step-Up in Basis Loophole Works," *Tech. rep.* [41]
- U.S. DEPARTMENT OF THE TREASURY (2022): *General Explanations of the Administration's Fiscal Year 2023 Revenue Proposals*, Department of the Treasury. [1]
- U.S. OFFICE OF MANAGEMENT AND BUDGET (2022): *Budget of the US Government, Fiscal Year 2023*, US Government Publishing Office. [1]
- VAN BINSBERGEN, JULES (2020): "Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility," *NBER Working Paper 27367*. [1]
- VICKREY, WILLIAM (1939): "Averaging of Income for Income-Tax Purposes," *Journal of Political Economy*, 47 (3), 379–379. [36]
- VON SCHANZ, GEORG (1896): "Der Einkommensbegriff und die Einkommensteuergesetze," *FinanzArchiv / Public Finance Analysis*, 13 (1), 1–87. [1]
- WERNING, IVÁN (2007): "Optimal Fiscal Policy with Redistribution," *The Quarterly Journal of Economics*, 122 (3), 925–967. [12]
- WHALLEY, JOHN (1979): "Capital Gains Taxation And Interest Rate Changes: An Extension Of Paish's Argument," *National Tax Journal*, 32 (1), 87–91. [6]
- WOLFF, EDWARD (2022): "The Stock Market and the Evolution of Top Wealth Shares in the United States," *The Journal of Economic Inequality*, 1–14. [6]
- ZHU, XIAODONG (1992): "Optimal fiscal policy in a stochastic growth model," *Journal of Economic Theory*, 58 (2), 250–289. [12]
- ZUCMAN, GABRIEL (2024): "A blueprint for a coordinated minimum effective taxation standard for ultra-high-net-worth individuals," *Tech. rep.*, available at <https://gabriel-zucman.eu/files/report-g20.pdf>. [1, 3]