

Post-Income Taxation: Financing the Agentic State

Paper III of *Foundations of Agentic Economics*

Matthew Long
The YonedaAI Collaboration
YonedaAI Research Collective
Chicago, IL
`matthew@yonedaai.com`

August 2026

Abstract

Papers I and II of this series argued that the primitive object of an economy is the obligation rather than money, and that autonomous agents can produce, contract and settle as economic principals. This paper sets out the fiscal consequence and argues for it: taxes on labour and productive income should be abolished, and the revenue replaced primarily by a broad tax on final consumption, with scarcity rents and the return on public capital as supplements rather than as the main instrument.

The case for abolition rests on two things this paper can show rather than assert. Income taxation is collected from wage earners with a completeness it never achieves against income that can be retained inside an entity, deferred, or converted into capital gain, because the third-party withholding that makes labour income administrable also leaves its holders no margin of avoidance. That is a claim about compliance and avoidance opportunity, not about economic incidence, which depends on price adjustments this paper does not model. And under agentic production the base itself becomes ambiguous, since an agent netting obligations against counterparties may never realise anything a twentieth-century code would recognise as receipts.

We then compute the replacement, honestly and per country, and report the rate the arithmetic actually requires rather than the rate an advocate would prefer. Assuming coverage of 96 per cent of the household consumption base, which is a stipulated ceiling rather than a measurement and which we take from New Zealand's revenue ratio on a wider base that includes government consumption, abolishing all taxes on labour income, corporate profit and social insurance in the United States requires a consumption tax of 37.2 per cent tax-exclusive, equivalently 27.1 per cent tax-inclusive. At the OECD-average coverage of 58 per cent it is 61.7 and 38.1. Adding a universal rebate on subsistence consumption raises

the first pair to 46.6 and 31.8 per cent. These are arithmetic requirements, not forecasts; no behavioural response is in them.

Denmark cannot do this. Replacing its productive-income gap together with all existing goods-and-services taxation means raising 41.35 per cent of GDP from a household base that, net of the taxes being replaced and at the same 96 per cent coverage, is 29.94 per cent of GDP: a rate of 138.1 per cent tax-exclusive, or 113.1 at Denmark's own current coverage. We print the failure rather than burying it. Norway and Singapore sit between. Singapore could abolish both income taxes immediately from investment returns with a consumption tax rise of a third of a percentage point, but only by spending the whole of its expected real return rather than the half its constitution permits, which converts a permanent endowment into a decaying one and reopens the gap within a generation.

The public capital endowment, which earlier framings of this argument treat as a precondition, is better understood as a rate reduction, and a bounded one. Abolition does not wait on it: a fund of one times GDP drawn at 4 per cent lowers the required American rate from 37.2 to 30.7 per cent tax-exclusive, and two times GDP to 24.2. But a fund that is drawn on while it accumulates converges rather than compounding, and at a 5 per cent contribution, a 4 per cent draw and a 2.5 per cent growth-adjusted return it converges to 3.33 times GDP. The consumption rate therefore falls toward a floor of 23.6 per cent tax-exclusive and never reaches zero. The endowment buys a lower rate, not the end of consumption taxation.

Keywords: consumption taxation, income tax abolition, optimal taxation, land value tax, sovereign wealth funds, agentic economics, public capital

JEL codes: H21, H23, H27, H60, Q32, O33

Contents

1	Introduction	5
2	Why Income Weakens as a Tax Base	8
2.1	The efficiency argument, and its limits	8
2.2	The administrative argument	9
2.3	Where the agentic surplus actually settles	11
2.4	What the argument does not claim	11
3	The Revenue Architecture	12
3.1	The decomposition	12
3.2	The public capital endowment	13
3.3	Two draw rules	14
3.4	Endowment equivalence	14
4	Method	15
4.1	What is being set to zero	15
4.2	The replacement gap	16
4.3	Yield models	16
4.4	Data and its limits	17
5	Country Tests	18
5.1	United States	18
5.2	Denmark	21
5.3	Norway	23
5.4	Singapore	26
5.5	United Arab Emirates	28
5.6	Cross-country summary	29
6	The Consumption Replacement	29
6.1	Rates, and the two ways of quoting them	29
6.2	What each country would have to charge	32
6.3	Reconciling with a coarser calculation	34
6.4	The rate at which a base stops behaving	34

6.5	Regressivity, and the instrument that answers it	36
7	What the Arithmetic Shows	37
7.1	The rent ceiling	37
7.2	What the Henry George Theorem actually delivers	37
7.3	The stock–flow confusion	39
7.4	The endowment lowers the rate toward a floor	40
7.5	What the fund can actually reach	40
7.6	Where the contribution comes from	42
7.7	Growth makes it harder	42
7.8	The shrinking externality base	43
8	Objections	44
8.1	Land valuation and the transition shock	44
8.2	Consumption taxes and distribution	45
8.3	The political economy of a very large state endowment	46
8.4	If taxes shrink, what anchors the currency?	46
8.5	Income survives because it is the broadest feasible base	47
9	A Transition Path	48
10	Limitations and Falsification	50
11	Conclusion	52
A	Derivations	53
A.1	The endowment ratio	53
A.2	Accumulation time	54
A.3	Land tax capitalisation	54
A.4	Consumption tax yield	54
B	Data Notes	55

1 Introduction

Tax systems are built out of whatever the state can see. The income tax, in the form that now finances most rich countries, became administratively possible only when a large share of economic activity started flowing through employers who kept written payrolls and could be compelled to report them. Kleven, Kreiner and Saez put the point sharply: modern governments can tax so much because firms act as fiscal intermediaries, and the paper trail generated inside a firm with many employees is difficult for any single participant to falsify [29]. The income tax is not a philosophical achievement. It is an accounting one.

That observation cuts both ways, and the second edge is the subject of this paper. If the income tax rests on an accounting fact about how production is organised, then a sufficiently large change in how production is organised puts the tax base itself in question. Paper I of this series argued, in a tradition running from Innes [20] and Knapp [30] through Graeber [17] and Mehrling [32], that the fundamental object in an economy is the obligation, and that money is the residue left over after obligations have been netted against one another, $M_t = \mathcal{R}(\mathcal{E}_t)$. Paper II argued that autonomous computational agents can occupy the vertices of that obligation graph as principals: discovering demand, negotiating, contracting, producing, subcontracting, and settling, with balance sheets of their own. That argument draws on the transaction-cost account of why firms exist at all [8, 60], on the agency-cost literature [24], and on a recent body of work treating AI systems as economic agents operating under institutional constraint [2, 18, 57]. Put the two together and something happens to the notion of taxable income. An agent that generates a large flow of real output while continuously netting obligations against counterparties may never realise anything that a twentieth-century tax code would recognise as receipts. Two economically identical arrangements, one settled in currency and one cleared inside a mutual credit network, can produce different tax liabilities. That is not tax avoidance. It is the tax base losing its referent.

That claim needs a boundary, and Paper II supplies two. Under current law in every jurisdiction its authors could identify, an autonomous agent cannot be a party to an enforceable contract; the obligor is an operator or a corporate wrapper. While that remains true, income accrues to an ordinary legal person with an ordinary balance sheet and stays perfectly administrable. Paper II also argues that capacity-based credit works soundest precisely where output is cheaply and objectively checkable, so the share of production financed against verified forward output is bounded by verifiability. The defensible version of the claim is therefore conditional: to the extent that autonomous actors transact by creating obligations against verified forward production rather than settling in currency, the realised-income base erodes for that share of activity, and that share is bounded by verifiability and by the availability of legal standing for agent obligations. Nothing in this paper requires the share to be large today.

The proposition this paper tests is the one stated in the series brief:

Axiom 1 (Productive neutrality). Production shall not constitute a taxable event.

In its pure form the axiom sets

$$\tau_{\text{labor}} = \tau_{\text{productive income}} = \tau_{\text{AI production}} = 0, \quad (1)$$

and relocates public revenue onto bases that remain observable and remain scarce:

$$T = T_{\text{land}} + T_{\text{resource rents}} + T_{\text{externalities}} + T_{\text{final consumption}} + R_{\text{public capital}}. \quad (2)$$

The theoretical pedigree of (1) is stronger than it may first appear. The Diamond–Mirrlees production efficiency theorem establishes that, under conditions that include the availability of a full set of commodity taxes and the taxation of pure profit, an optimal tax system should never distort the aggregate production plan [11, 12]. Taxes belong on final transactions, not on intermediate ones. Henry George had reached a related conclusion a century earlier by a different route: the return to land is an unearned increment arising from location and scarcity rather than from effort, and taxing it costs society nothing in foregone production because the supply of land does not respond [15]. The Mirrlees Review, writing in 2011 with no interest whatsoever in autonomous agents, recommended a land value tax on precisely these grounds and called it the closest thing to a neutral tax that exists [36].

The case for abolition

This paper argues for abolition and then prices it. Both halves matter, and the second is where most advocacy of this kind fails.

The normative case does not depend on the claim that income taxation was designed to impoverish anyone. It was not, and Section 2.4 says so plainly. The case depends instead on a structural asymmetry that the administrative literature establishes without intending to. Kleven, Kreiner and Saez show that income taxation works because third-party reporting inside firms makes wages nearly impossible to conceal [29]. Read as a description of where the tax is actually collected, that mechanism has an uncomfortable corollary. Withholding applies with near-perfect completeness to people whose entire income arrives as wages, and with much less completeness to income that can be retained inside an entity, deferred, converted into capital gain, or realised at a time of the holder’s choosing. The taxpayer with only labour has no such margins.

The claim needs stating precisely, because it is easy to overstate into something the evidence does not support. This is an argument about compliance and avoidance opportunity, not about economic incidence. Kleven, Kreiner and Saez is an administrability result; where the burden finally settles after wages, prices and asset values adjust is a general equilibrium question that

neither their paper nor ours addresses. What can be said is narrower and still uncomfortable: the completeness with which the statutory tax is collected varies systematically with the composition of a household's income, and it is most complete for households whose income is entirely wages. A tax whose collection rate is determined by the shape of the enforcement mechanism rather than by the schedule is a blunt instrument for distributive justice even when the schedule is graduated.

The second half of the case is the one this series contributes. Under agentic production the object being reported becomes ambiguous even where reporting is perfect, for the reasons set out in Section 2. Together these give a direction of travel: away from taxing the act of producing, and toward taxing what is finally consumed and what is genuinely scarce.

What the arithmetic has to answer

The literature on replacing income taxation has a long history of elegant arguments and a much thinner history of arithmetic at national scale, and the reason is that the arithmetic tends to be unkind. So we take five countries, decompose their actual revenue, compute what would have to be replaced, and report the rates that follow. The countries span the range rather than making a case. Denmark and the United States are the hard cases. Norway and Singapore are the closest things to working demonstrations. The United Arab Emirates is the closest realised approximation to $\tau_{\text{labor}} = 0$ anywhere on earth, and both why it works and what it is resting on repay attention.

One methodological commitment governs everything that follows. Where a rate is quoted we give it both tax-exclusive and tax-inclusive, because the difference between those two conventions is the single most common source of confusion in this debate and the gap between them widens as rates rise. A 30 per cent tax-exclusive rate is 23 per cent tax-inclusive; the same tax described two ways sounds like two different policies.

Three results can be stated in advance, since the rest of the paper is essentially their derivation.

First, the replacement gaps are very large. Taxes on labour income, business profits and social insurance account for 72 per cent of all tax revenue in the United States and 66 per cent in Denmark. In levels these are 18.5 and 28.9 per cent of GDP respectively. There is no plausible configuration of land, carbon and consumption taxation that closes a gap of that size, and we show this by computing the implied rates and comparing them to the highest rates ever observed anywhere.

Second, the only pillar in (2) with the arithmetic capacity to close such a gap is the public capital endowment, and its capacity is a function of an accumulated stock rather than a policy rate. This makes the proposal a statement about national saving over multiple generations.

We formalise the requirement as an equivalence between the revenue share to be replaced and the endowment-to-GDP ratio needed to replace it, and we show that the required ratios sit between 1.4 and 12.4 times GDP depending on the country and the draw rule. For the United States the upper figure corresponds to an endowment larger than the entire net worth of the American household sector.

Third, the transition contains a genuine paradox that the proposal in its usual form does not acknowledge. A state that intends to abolish production taxes by accumulating capital must run large surpluses for decades, and those surpluses have to be raised from the tax system it currently has. The endowment is financed by the income tax until the moment the income tax is no longer needed. Anyone advocating this path is advocating higher taxes on production for fifty to a hundred years in exchange for lower taxes afterwards.

Section 2 sets out the argument against income as a base in an agentic economy, and marks carefully where the argument stops. Section 3 develops the revenue architecture and the endowment recursion, and proves the equivalence result that organises the quantitative work. Section 4 states the method and its limitations. Section 5 contains the country tests. Section 7 draws out what the arithmetic shows, including the failures. Section 8 takes up four serious objections. Sections 9 and 10 discuss transition and falsification.

2 Why Income Weakens as a Tax Base

2.1 The efficiency argument, and its limits

The modern theory of taxation begins with Ramsey’s problem of raising a fixed revenue at least welfare cost [48] and with Mirrlees’s reformulation of the redistributive problem under asymmetric information [37]. The particular result that matters here arrived in 1971.

The production efficiency theorem is the natural formal ancestor of Axiom 1. Diamond and Mirrlees showed that if a government can tax all final consumption goods and can capture pure profit, then the optimum involves no distortion of the aggregate production plan: intermediate transactions should go untaxed, and the marginal rate of transformation should equal the producer price ratio [11]. Under those conditions a tax on business inputs or on the return to reproducible capital is dominated by a tax on final transactions raising the same revenue.

The theorem has well-understood conditions, and honesty requires stating them. It assumes constant returns or full profit taxation, an unrestricted set of commodity taxes, and no untaxed sector. Where any of these fails, production taxes can re-enter the optimum. The long-run zero capital income tax results of Judd [26] and Chamley [7] pointed in the same direction and were for decades treated as settled. Straub and Werning have since shown that the Chamley–Judd conclusion depends on convergence properties that fail for a range of standard preferences,

and that the optimal long-run capital tax can be strictly positive [55]. Atkinson and Stiglitz, meanwhile, established that when preferences are weakly separable between consumption and leisure and a nonlinear income tax is available, *differential* commodity taxation adds nothing [4]. That result is often misquoted in debates of this kind, so precision about what it does and does not say matters here. The theorem is this: under weak separability between consumption and leisure, and with a nonlinear income tax available, *differential* commodity taxation is redundant. That is a statement about varying rates across goods. It says nothing about whether to tax consumption or income.

A separate and much simpler proposition is often run together with it, including by us in an earlier draft of this section. A flat consumption tax and a flat labour income tax are close to equivalent over a lifetime, absent initial wealth and bequests. That is a budget-constraint identity, not the theorem: it follows from the fact that lifetime labour income is either consumed or bequeathed, and it requires none of Atkinson and Stiglitz’s separability machinery. The bite of Atkinson–Stiglitz for this paper is therefore narrower and sharper than a general warning. Efficiency survives the move from a nonlinear income tax to a uniform consumption tax. What does not survive is the capacity to set marginal rates that vary with ability, which is to say the redistributive instrument itself. We return to that in Section 8.

None of this is fatal to the proposition, but it means the efficiency case cannot carry the argument alone. What the theory establishes is a presumption in favour of taxing final transactions and pure rents rather than intermediate production, together with the observation that the presumption is conditional. The interesting question is whether the conditions are becoming more or less satisfied over time.

2.2 The administrative argument

The stronger claim in an agentic setting is administrative rather than allocative. Kleven, Kreiner and Saez model the firm as a fiscal intermediary and show that third-party reporting sustains high tax compliance because collusion between many employees is hard to organise and easy to detect [29]. Gordon and Li make a related point about the financial sector: tax systems in developing countries look the way they do because the observable margin is the boundary of the formal financial system [16]. Slemrod’s framing of an “optimal tax system” rather than optimal tax rates makes administration a first-class object of design [52].

An agentic economy changes the information environment in two directions at once, and it matters that they point opposite ways.

In one direction, verification improves dramatically. Paper IV of this series specifies an identity and contracting layer in which each participant holds a cryptographically verifiable economic identity and contracts are machine-readable objects with defined verification and settlement conditions. If that infrastructure exists, third-party reporting is no longer an

approximation extracted from reluctant employers; it is a structural property of the settlement layer. On this reading the administrative case for the income tax gets *stronger*, not weaker, and a referee is entitled to press the point.

In the other direction, the object being reported becomes ambiguous. This is the part that matters, and it follows directly from Paper I. Consider three arrangements. In the first, an agent sells services for currency and receives \$50 million. In the second, the same agent supplies the same services under a standing contract and receives compute, energy, data access and physical logistics of equal value, with no currency changing hands. In the third, the agent operates inside a clearing network where its obligations to counterparties are netted continuously against theirs, and its net settlement position over the year is close to zero despite gross throughput of hundreds of millions.

All three arrangements represent the same real production. Only the first is unambiguously income under existing law. The second is barter and in principle taxable, though in practice the valuation problem is severe when the goods exchanged are non-standard and the counterparties are related. The third is not obviously anything at all: the netting is exactly the operation Paper I describes as the passage from the obligation graph $\mathcal{E}_t$ to the residual $\mathcal{R}(\mathcal{E}_t)$, and there is no principled point in that process at which one is compelled to say that income has been realised.

The realisation problem is not new. Tax law has struggled with it since the invention of the corporation, and every jurisdiction has an elaborate apparatus of constructive receipt, transfer pricing and controlled foreign corporation rules to manage it. What changes under agentic production is the volume and the speed. The devices that currently require a tax lawyer and a year of planning become the default behaviour of a clearing system optimising for settlement efficiency. When netting is a continuous background process rather than a deliberate structuring choice, the boundary between ordinary operations and avoidance stops being administrable, because there is no intent to detect.

One qualification, since the argument is easy to push too far. Perfect visibility into transactions does not by itself solve the problem, because the residual difficulty is valuation rather than observation. When an agent exchanges compute for logistics with a related counterparty, a tax authority that can see the transaction in full detail still has to decide what it was worth, and there may be no arm's-length comparable because the bundle is unique. This is the transfer pricing problem, and it is not obviously easier when the parties are software. What agentic settlement changes is the volume of transactions to which the problem applies, which is a reason to expect the base to become harder to administer, not a demonstration that it becomes impossible.

2.3 Where the agentic surplus actually settles

The preceding argument is conditional and negative: it says a base may erode. There is a second argument available which is unconditional and positive, and it does more work for this paper than the erosion claim does. It comes from Paper II and we set it out in full because it changes the status of the revenue architecture in (2).

Suppose agent capability is replicable at negligible marginal cost and entry into agent operation is free. Then competition drives the price of agentic output toward marginal operating cost, and the pure return to autonomous intelligence goes to zero. The surplus does not vanish; it moves. It accrues to whatever the agents need that cannot be replicated on demand: energy generation and transmission, sited data-centre capital, land, spectrum, proprietary non-replicable data, and non-transferable reputation. Every item on that list except the last two is already a pillar of (2).

This reframes the whole proposal. Taxing land, energy, resources and spectrum is not a philosophical preference for rent over production, nor a second-best consolation once income becomes unadministrable. It is a claim about where the returns to an automated economy physically come to rest. If Paper II's competitive condition holds, the scarcity bases are not a substitute for taxing the agentic surplus. They *are* the agentic surplus. That argument survives even if the legal standing gap never closes and income remains perfectly measurable, which is why we lead the theoretical case with it and treat the erosion of realised income as secondary.

The condition carries its own warning, and it is a serious one. The argument assumes model weights are freely replicable. Frontier models are not distributed that way: access is controlled by contract and by secrecy, and where that holds the provider captures the rent instead of the inelastic inputs. Provider rent is mobile, intangible, and locatable almost anywhere for tax purposes, which is to say it has precisely the properties that land does not. A pillar architecture built on immobile bases has no instrument for it. This is a live weakness rather than a hypothetical one, and Section 10 returns to it.

2.4 What the argument does not claim

Three disclaimers, because the proposition is easy to overstate and the source discussion that motivated this series had to be corrected on exactly this point.

We do not claim that income taxation was designed to impoverish anyone. Income taxes have historically funded public services, redistributed resources and shaped behaviour, and the historical record on their purposes is not seriously in dispute. The argument here is about distortion and about the eroding meaning of the base, which is a different and much narrower claim.

We do not claim that the income tax is currently failing. It is not. Compliance rates

in the countries studied here remain high, and the tax raises very large sums at moderate administrative cost. The claim is conditional on a transformation of production that has not happened yet and may not happen at the scale assumed.

We do not claim that a zero rate is achievable everywhere. Section 5 shows that it is not, and Article VII of the constitutional framework in Paper V is drafted accordingly. It schedules the taxation of labour and productive income for abolition rather than declaring it abolished, makes final consumption the primary base with rents, royalties and public-capital returns as supplements, and imposes an annual duty to disclose what share of revenue is still raised from production and on what trajectory. It also provides an emergency departure clause, which must be announced, is reviewable, and is limited in time. A schedule with a disclosure duty and a bounded escape hatch is a different instrument from an aspiration, and it is the one this paper's arithmetic supports.

3 The Revenue Architecture

3.1 The decomposition

Equation (2) partitions revenue into five pillars, each corresponding to a base that survives the argument of the previous section because it is defined by scarcity or by the absorption of real resources rather than by an accounting event.

The pillars are not equal partners and this paper does not treat them as such. Final consumption taxation is the primary instrument, for the plain reason that it is the only one of the five with a base large enough to carry the load. It applies where claims are converted into the absorption of scarce real goods and services, which is the point Paper I identifies as settlement into consumption rather than as the creation of an obligation. Section 6 computes what it must raise and at what rate.

The remaining pillars are supplements, and saying so is a departure from how this argument is usually made. Land value taxation falls on the unimproved value of location; its supply is fixed and a tax on its rent is capitalised into a lower purchase price without altering the quantity supplied [56]. Resource rents cover petroleum, minerals, spectrum, water rights, rights of way, and any other scarce entitlement whose value derives from exclusivity rather than from production. externality charges price measurable harm at its social cost, following the Pigouvian logic formalised by Sandmo [51]. Each of these is a good tax on its merits. None is a large one: Section 7 shows that no country has ever raised more than 1.09 per cent of GDP from a pure land tax or more than 0.44 per cent from carbon. Treating them as the main replacement for a base worth eighteen to twenty-nine points of GDP is an evasion dressed as a programme.

The fifth pillar, the return on publicly held capital, is different in kind and is developed below.

At the level of an individual participant the first four pillars give the per-agent form

$$T_i = \tau_C C_i + \tau_L L_i + \tau_R R_i + \tau_E E_i, \quad (3)$$

with C_i final consumption, L_i land value held, R_i extraction of scarce natural resources, and E_i measurable externality cost imposed. The tax event moves from production to consumption:

$$\text{claim} \rightarrow \text{final consumption} \rightarrow \text{tax event}, \quad \$1\text{M earned} \neq \$1\text{M taxable}. \quad (4)$$

Remark 1 (Notation). Papers I and II use L_i for the liabilities of agent i and $C_i(t)$ for its economic capacity. Within this paper L_i denotes land value and C_i final consumption, as in (3). The scope is local to the fiscal notation and we do not use the balance-sheet meanings here.

3.2 The public capital endowment

The fifth pillar is a stock rather than a rate, which is what makes it both the most powerful element of the architecture and the most easily misdescribed. Let $W_G(t)$ denote a publicly owned, diversified capital endowment. It evolves as

$$W_G(t+1) = W_G(t) + I_G(t) + R_G(t) - G_t, \quad R_G(t) = r_t W_G(t), \quad (5)$$

where $I_G(t)$ is new contribution from the public purse, $R_G(t)$ the return earned, and G_t public expenditure financed from the fund in period t .

The motivating idea is that a rising share of output produced by autonomous capital can accrue to collectively owned assets without the state owning or directing the productive economy. A sovereign fund holding a diversified global portfolio is a passive claimant on production, not a planner of it. This distinction is what allows the arrangement to sit inside a framework that is neither conventionally capitalist nor socialist, and it is the same distinction that separates Norway's Government Pension Fund Global from state ownership of Norwegian industry.

The fiscal condition the architecture must satisfy is

$$R_G + R_{\text{land}} + R_{\text{resources}} + T_{\text{externalities}} + T_{\text{consumption}} \geq G, \quad (6)$$

with productive-income taxation approaching zero.

3.3 Two draw rules

Everything quantitative in this paper turns on a distinction that is usually left implicit in discussions of sovereign funds, so we make it explicit.

Let Y_t be GDP growing at real rate g , let the fund earn real return r , and write $w_t = W_G(t)/Y_t$ for the endowment expressed as a multiple of GDP. Suppose the state withdraws a constant share ϕ of the fund each year. Then in continuous time

$$\dot{w} = (r - g - \phi) w, \quad (7)$$

so the fund's ratio to GDP is stationary if and only if $\phi = r - g$.

This gives two natural policies. Under what we will call the *consumption rule* the state spends the whole real return, $\phi = r$. The fund then retains its real value but shrinks steadily relative to the economy, at rate g . Under the *stewardship rule* the state spends $\phi = r - g$, and the fund holds its position as a share of national output indefinitely. The two rules differ by a factor of $r/(r - g)$ in the endowment required to finance a given revenue share, and at plausible parameters that factor is not small. With $r = 4.5\%$ and $g = 2\%$ it is 1.8.

The distinction is not academic. Singapore's Net Investment Returns framework permits the government to spend up to half of the expected long-term real return on the net assets managed by GIC, MAS and Temasek. If the expected real return is around 4.5%, then half of it is 2.25%, which is approximately $r - g$ for a mature economy growing at a little over two per cent. Singapore's constitutional rule is, to a good approximation, the stewardship rule. Norway's fiscal guideline, which targets a draw of three per cent against an expected real return of the same order, is closer to the consumption rule, and the fund's continued growth relative to mainland GDP has come from new petroleum contributions rather than from retained returns.

3.4 Endowment equivalence

We can now state the result that organises the country tests.

Proposition 1 (Endowment equivalence). *Let $\gamma = T^{prod}/Y$ be the share of GDP currently raised from taxes on production that the reform would set to zero, and suppose the entire gap is to be financed from the return on a public capital endowment under a draw rule with sustainable withdrawal rate ρ . Then the required stationary endowment is*

$$w^* = \frac{W_G}{Y} = \frac{\gamma}{\rho}, \quad (8)$$

with $\rho = r$ under the consumption rule and $\rho = r - g$ under the stewardship rule.

Proof. Under a constant withdrawal share ϕ , revenue from the fund is ϕW_G , and the requirement

is $\phi W_G = \gamma Y$, hence $W_G/Y = \gamma/\phi$. Stationarity of W_G/Y requires $\phi \leq r - g$ with equality at the maximal sustainable draw; setting $\phi = \rho$ gives (8). Under the consumption rule the real value of the fund is preserved, $\phi = r$, and (8) holds at each date but w declines at rate g , so the revenue share γ is sustained only transiently. $\square$

The content of Proposition 1 is the multiplier $1/\rho$. At a sustainable draw of 2.5 per cent, every percentage point of GDP of income tax to be abolished requires an endowment of 40 per cent of GDP. That factor of forty is the single most important number in this paper, and it is what converts a proposal about tax rates into a proposal about several generations of national saving.

Corollary 1 (Accumulation time). *Starting from $w(0) = 0$ and contributing a constant share s of GDP each year to a fund earning real return r in an economy growing at g , with no withdrawals, the endowment ratio follows $\dot{w} = s + (r - g)w$, so*

$$w(t) = \frac{s}{r - g} \left(e^{(r-g)t} - 1 \right), \quad t^* = \frac{1}{r - g} \ln \left(1 + \frac{w^*(r - g)}{s} \right). \quad (9)$$

Corollary 1 is what makes the transition tractable and, as it turns out, discouraging. Notice what (9) does not contain: any term that gets better if artificial intelligence raises productivity growth. Faster growth raises g , which lowers $r - g$, which lowers the sustainable draw and *raises* the required endowment under the stewardship rule. An AI productivity boom makes the accumulation target harder to hit relative to GDP even as it makes the absolute sums easier to generate. We return to this in Section 7.

4 Method

4.1 What is being set to zero

We define the productive-income tax aggregate as the sum of personal income tax on labour and business income, corporate income tax, social security contributions, and payroll taxes. All four are levied on production or on the act of employing labour, and all four are within the scope of (1). Payroll taxes are negligible in the United States and Norway, at 0.03 and 0.05 per cent of GDP, and zero in Denmark, but we include them for consistency and note where they move a total. Social security contributions are included because they are a payroll tax in economic substance whatever their legal characterisation, and because excluding them would understate the gap by six percentage points of GDP in the United States.

The choice is contestable and we test it rather than defend it. Social contributions partly purchase an entitlement, and to the extent that a worker perceives a link between contributions paid and benefits received, the contribution is deferred compensation and distorts labour supply

less than a pure tax. Against that, the link is weak in most systems and close to absent at the margin in several. Table 5 therefore reports a robustness row for the United States in which social contributions survive the reform: the gap falls from 18.51 to 12.48 per cent of GDP and the required stewardship endowment from 7.40 to 4.99 times GDP. The American verdict does not change, though the magnitude of its failure does. Denmark and Singapore are unaffected, since Danish social insurance is financed through the income tax and Singapore’s Central Provident Fund is a mandated individual savings scheme rather than a tax.

We exclude from the aggregate: taxes on goods and services, recurring taxes on immovable property, resource royalties and production taxes on extraction, and non-tax revenue. These are either already among the five pillars or can be mapped onto them.

Two categories sit awkwardly and we flag them rather than resolve them. Capital gains taxation is neither clearly a tax on production nor clearly a tax on rent, since realised gains mix genuine saving returns with capitalised land and monopoly rents. We leave it inside the personal income tax aggregate, which makes our gap estimates slightly conservative in the direction of being too large. Corporate income tax on resource extraction, which in Norway includes a marginal petroleum rate of 78 per cent, is formally a profits tax but economically a resource rent charge; we discuss the Norwegian case separately for this reason.

4.2 The replacement gap

For each country we compute

$$\gamma = \frac{T_{\text{PIT}} + T_{\text{CIT}} + T_{\text{SSC}}}{Y}, \quad (10)$$

expressed in percentage points of GDP. This is the revenue that must be found elsewhere if the reform holds total revenue constant.

Holding revenue constant is a choice and it deserves defending. One could instead ask what happens to the deficit, or assume expenditure falls. We hold revenue constant because the alternative conflates two arguments: a claim about the structure of taxation and a claim about the size of the state. The proposition under test says nothing about how large government should be, and testing it jointly with an implicit spending cut would be a way of smuggling in a conclusion. Where a country currently runs a deficit, as the United States does substantially, our test therefore understates the true financing requirement, and we say so where it matters.

4.3 Yield models

For consumption taxation we use the standard decomposition in which revenue equals the statutory rate times the base times the VAT revenue ratio, the OECD’s measure of how much of a theoretically comprehensive base is actually reached after exemptions, reduced rates and non-compliance [43]. Writing c for household final consumption as a share of GDP and κ for

the share of that base the tax actually reaches, each percentage point of consumption tax yields approximately $c\kappa$ percentage points of GDP. We report results at two coverage levels: 0.96, the highest ratio achieved anywhere, and 0.58, the OECD average. Note that κ is defined here on the household base and is not the OECD’s VAT Revenue Ratio, whose denominator includes government consumption; Section 6 sets out why the two cannot be interchanged and in what sense we borrow the 0.96.

For land we invert the question. Rather than assert a land value stock, we compute the stock that would be required: to raise γ percentage points of GDP at an annual rate τ_L on land value V_L requires $V_L/Y = \gamma/\tau_L$. We then compare the requirement to the highest recurring property tax yields actually observed anywhere, which is the relevant empirical ceiling and is considerably more informative than a debate about land value estimates.

For externality charges we take physical emissions and multiply by a price, using the EPA central social cost estimate as the upper anchor [59]. This is generous, since it prices all emissions at the full social cost with no exemptions and no international competitiveness carve-outs, neither of which any real carbon pricing system has managed.

For the endowment we apply Proposition 1 at three draw rates: 5 per cent and 4 per cent, bracketing the range specified in the series brief and corresponding to aggressive and moderate versions of the consumption rule; and 2.5 per cent, corresponding to the stewardship rule at $r = 4.5\%$ and $g = 2\%$.

4.4 Data and its limits

Revenue composition comes from OECD Revenue Statistics [45] and, for Singapore, from the Asia-Pacific companion volume [46] together with the Singapore Ministry of Finance’s own revenue and expenditure analysis [33] and the tax authority’s annual report [19]. Norwegian fund and fiscal rule material comes from the IMF Article IV report [22] and the Ministry of Finance long-term perspectives white paper [41]. UAE aggregates come from the IMF Article IV [23].

Several figures used here are approximations and we mark them as such at the point of use rather than in a footnote that readers will skip. Vintages differ, and we do not force them onto a common year. All three OECD countries are reported for 2023, the most recent year for which the 2025 edition publishes a full category breakdown for every country in the group. Singapore is reported for fiscal year 2025 and the UAE for 2023, neither of which can be aligned to a calendar-year OECD basis; we flag the mismatch rather than interpolate. Ratios are taken from the OECD edition throughout and are not mixed with nationally computed ratios, since the two use different GDP vintages. For the United States the difference is not trivial: the OECD ratios rest on a GDP of \$27.72 trillion for 2023 while the current BEA vintage is \$27.81 trillion, so a ratio computed from one and a level from the other will not reconcile.

A trap in the Norwegian series. A naive reading of the OECD compilation records Norwegian personal income tax at approximately 19.5 per cent of GDP. Combined with corporate income tax at 9.8, social contributions at 9.5 and value added tax at 7.5, the components then sum to 46.3 per cent of GDP against a published total of 41.6, which is impossible. The cause is that 19.5 is category 1000, which is personal *and* corporate income taxes combined. Because Norway’s category 1300 is zero, 1000 is exactly 1100 plus 1200, so reading it as a personal income tax line and then listing corporate tax separately counts the corporate tax twice and produces the entire 4.7 point overshoot. The correct Norwegian personal income tax is 9.60 per cent of GDP in 2023 and 10.00 in 2024.

The trap is worth naming because it is the kind of error that survives review easily. Nothing about 19.5 looks wrong for a Nordic country, the components are individually plausible, and only the sum betrays it. Anyone reproducing this exercise should check that components sum to the published total before doing anything else, and should treat category 1000 as an aggregate rather than as a personal income tax line.

A figure we could not verify. The mainland Norwegian tax ratio of 42.8 per cent quoted in several secondary sources is not reproducible from primary data by any construction we could identify. Total tax over mainland GDP gives 53.0 per cent; total tax net of petroleum taxes over mainland GDP gives 41.3. We use the latter, state the definition wherever it appears, and do not use 42.8.

5 Country Tests

5.1 United States

The United States has the lowest total tax take in our group and the most income-dependent structure. Table 1 decomposes the 25.6 per cent of GDP collected in 2023.

The gap is 18.51 per cent of GDP, or roughly \$5.1 trillion at 2023 output of \$27.72 trillion on the OECD vintage. This understates the financing problem twice over. General government expenditure in the United States was 37.7 per cent of GDP in 2024 against revenue of about 25.6, so holding revenue constant holds a deficit of roughly twelve points of GDP constant with it. And the country’s total tax take is already the third-lowest in the OECD, against an average of 33.7 per cent, so there is no fat in the existing system to redirect.

Consumption. American household final consumption is unusually high, at roughly 68 per cent of GDP, which is the one favourable feature of the United States case. Consolidating the existing state and local sales taxes into a single national consumption tax and using it to replace the whole 22.76 per cent of GDP requires, from Table 7, 37.2 per cent tax-exclusive on a base

Table 1: United States: revenue composition, 2023. Both columns are published OECD figures, not derived from one another; the left column is a share of total tax revenue and the right a share of GDP [45]. Components sum to 25.67 against a published total of 25.63, and the share column to 100.1 per cent, because of rounding. The property line is OECD category 4000, all taxes on property; the narrower category 4100, recurring taxes on immovable property, is 2.67 per cent of GDP and is the figure used in the land discussion of Section 5.

Base	% of total tax revenue	% of GDP
Personal income tax	40.0%	10.25
Social security contributions	23.5%	6.03
Taxes on property (4000)	11.3%	2.91
Goods and services	16.6%	4.25
Corporate income tax	8.6%	2.20
Payroll	0.1%	0.03
Total	100.0%	25.63
Productive-income aggregate	72.2%	18.51

with New Zealand’s coverage and 61.7 per cent at the OECD-average coverage of 0.58. The tax-inclusive equivalents are 27.1 and 38.1. The highest standard rate in force anywhere is 27 per cent, in Hungary, followed by Finland at 25.5 and a cluster at 25; no country outside Europe reaches 23. The United States would need to exceed the world maximum while simultaneously achieving the best coverage ever recorded, in a country that has never operated a value added tax at all. Both figures are arithmetic requirements rather than forecasts: the behavioural response is not in them, for the reasons in Section 6.4.

A note on the coverage parameter, since it is easy to misuse. The OECD’s VAT Revenue Ratio averaged 0.58 across member countries in 2022 and stood at 0.66 for Denmark [43], but its denominator is total final consumption including government, not household consumption, so the two are not interchangeable. We therefore use an explicit coverage share of household consumption rather than importing a VRR into a base it was not constructed for. The OECD has published no VRR after 2022.

Land. The empirical ceiling here is stark, and the American land stock is smaller than the Georgist case usually assumes. Recurring taxes on immovable property already raise 2.67 per cent of GDP, the second-highest in the OECD after the United Kingdom, and that figure taxes buildings as well as land.

On the stock itself, Larson’s estimate for the Bureau of Economic Analysis puts all American land at \$23.0 trillion in 2009, which was 1.59 times GDP [31]; Davis and Heathcote’s residential land series is consistent with that order of magnitude [9]; and a residual calculation from the Federal Reserve’s financial accounts against BEA structures gives private land of roughly \$28 trillion in 2024, or 0.96 times GDP [6]. The BEA publishes no official land balance sheet, which

is why the range is wide. For comparison, OECD national accounts put land at 2.4 to 2.8 times GDP in France, the United Kingdom and Japan, and 3.8 in Australia [44], so the United States sits at the bottom of the advanced-economy distribution rather than the middle.

Take one to 1.6 times GDP. At a four per cent market yield that is an aggregate land rent of 4.0 to 6.4 per cent of GDP, and that rent is the absolute ceiling on what any land tax can raise, since it is the entire flow the asset produces. Against a gap of 18.51 per cent of GDP, complete confiscation of every dollar of American land rent covers between 22 and 35 per cent of the requirement. A tax on assessed capital value does worse still, for the reason set out in Section 7.2: at a 3 per cent rate it captures only 43 per cent of the rent. Reaching 18.51 per cent of GDP that way would require a pre-tax land stock of 10.8 times GDP, roughly seven times the largest published American estimate.

Externalities. The United States emitted roughly 4.8 to 4.9 billion tonnes of fossil carbon dioxide in 2024, the range reflecting genuine disagreement between the energy-related and territorial fossil-plus-industry definitions rather than uncertainty within either [21]. Priced at the EPA central social cost figure of \$190 per tonne, that is about \$0.93 trillion, or 3.3 per cent of GDP [59].

Three qualifications, all of which cut the same way. The \$190 figure is specific to a 2.0 per cent near-term Ramsey discount rate for 2020 emissions in 2020 dollars; at 2.5 per cent the same report gives \$120 and at 1.5 per cent \$340, and estimates across the wider modelling literature vary by roughly an order of magnitude [39]. The figure is also no longer operative United States regulatory policy: Executive Order 14154 of January 2025 disbanded the interagency working group and withdrew its guidance, so we are pricing at a published technical estimate rather than at a rate any government currently applies. And the calculation assumes every tonne is taxed at full social cost with no exemption and no competitiveness carve-out.

Against that theoretical 3.36 per cent stands the observed record. No country has ever raised more than one per cent of GDP from carbon taxation; the verified maximum is Norway at 0.44 per cent of GDP in 2024, counting both the general and the offshore petroleum carbon tax. Global carbon pricing revenue of roughly \$102 billion is about 0.09 per cent of world GDP [61]. The theoretical American maximum is therefore about seven and a half times the best performance any country has achieved. It is also, by construction, a declining base: a Pigouvian tax set at the social cost is meant to reduce the activity, and a carbon price that works eliminates its own revenue.

Endowment. Applying Proposition 1, financing 18.51 per cent of GDP requires an endowment of 3.70 times GDP at a 5 per cent draw, 4.63 times at 4 per cent, and 7.40 times under the stewardship rule at 2.5 per cent. In 2023 dollars those are \$103 trillion, \$128 trillion and \$205

trillion. Household and nonprofit net worth in the United States was \$183.0 trillion in the first quarter of 2026 [6]. The moderate case requires a public endowment equal to about 70 per cent of everything American households own; the stewardship case requires 112 per cent of it, which is to say more than the entire net worth of the household sector. Corollary 1 puts accumulation at 93 years if the state saves 2 per cent of GDP annually and 62 years at 5 per cent, assuming $r = 4.5\%$ and $g = 2\%$.

Verdict. Abolition is arithmetically available to the United States and the price is a consumption tax of 37.2 per cent tax-exclusive, 27.1 per cent tax-inclusive, on a base with New Zealand’s coverage; 46.6 per cent with a universal rebate. The word arithmetically is doing real work in that sentence. These are the rates that balance the identity, not rates any tax authority has been shown to be able to collect; the behavioural response is not in them, and Section 6.4 explains why at these levels that omission is not a technicality. That is a rate no country has levied and a base no country except New Zealand has achieved, and both would have to be delivered at once by a state that has never operated a national consumption tax. It is not impossible and we do not describe it as impossible. It is unprecedented, which is a different claim and a defensible one.

What the scarcity pillars contribute here is modest and should be stated as such. Full confiscation of American land rent covers 22 to 35 per cent of the gap, carbon at full social cost 3.3 points, and both together leave the great majority of the requirement to consumption. A land tax and a carbon price are worth having on their own merits. They are not an alternative to the consumption rate in Table 7; they lower it by a few points.

5.2 Denmark

Denmark is the hardest case in the world for this argument, which is why we include it. It has the highest tax take in the OECD, it funds a comprehensive welfare state almost entirely from income taxation, and it has essentially no resource rent.

The gap is 28.94 per cent of GDP, or 65.7 per cent of all Danish tax revenue. Denmark’s total take rose to 45.2 per cent of GDP in 2024, the highest in the OECD, but we run the analysis on 2023 so that all three OECD countries share a year.

One classification point matters for what follows. The Danish property value tax (*ejendomsværdiskat*) is recorded by both the OECD and Statistics Denmark under personal income tax, not under property, because it is levied on imputed income from owner-occupied housing. Anyone who moves it into the property line will overstate Danish property taxation by about half a point of GDP and understate the income tax by the same amount [53].

Table 2: Denmark: revenue composition, 2023, all cells as a percentage of GDP [45]. Social contributions are near zero because social insurance is financed through the income tax; the labour market contribution (*arbejdsmarkedetsbidrag*) is classified as personal income tax, which is why that line is so large. The residual is unallocated minor categories.

Base	% of GDP
Personal income tax	25.17
Corporate income tax	3.69
Social security contributions	0.08
Goods and services (of which VAT 9.14)	12.41
Taxes on property	1.72
Other and unallocated	0.97
Total	44.04
Productive-income aggregate	28.94

Consumption. Denmark already levies a 25 per cent VAT which raises 9.14 per cent of GDP against a household final consumption base of 43.6 per cent [62]. Replacing the productive-income gap together with the whole of goods-and-services taxation means raising 41.35 per cent of GDP from consumption. From Table 7 that is 138.1 per cent tax-exclusive at New Zealand’s coverage and 228.6 at the OECD average; the tax-inclusive equivalents are 58.0 and 69.6. At Denmark’s own current coverage of 0.66 the requirement is 113.1 per cent tax-exclusive.

Every one of those figures is far outside anything observed, and all of them ignore the behavioural response, which at such rates would not be marginal. A 100 per cent consumption tax is a 50 per cent tax-inclusive rate, at which cross-border shopping, informal exchange and outright evasion would collapse the base long before the revenue arrived. Denmark’s small size and open land border with Germany make this worse rather than better.

Land. Denmark runs the best natural experiment available for the Georgist proposition. The *grundskyld* is a genuine land value tax, levied on assessed land value separately from buildings, and since 2022 the business property charge (*dækningssafgift*) has also been assessed on land value, so the whole of Denmark’s recurring immovable property category is now land-based.

The yields are small and, more instructively, unstable. Grundskyld raised 1.09 per cent of GDP in 2023 and 0.77 in 2024; the land-based aggregate was 1.21 per cent falling to 0.89 [53]. The fall of roughly a quarter in a single year was not a collapse in land values but a statutory reform, which moved assessments onto a market-based footing and cut municipal rates from a 16.0–34.0 per mille range to 3.10–17.70. The 2024 and 2025 bills are provisional pending final valuations. A tax base whose yield can be cut by a quarter in one year by ordinary legislation is not obviously the immovable foundation the Georgist argument imagines, whatever the theory says about the immobility of land itself.

The ceiling binds hard here. Danish land is not especially valuable relative to output: OECD national accounts put household land at about 0.53 times GDP [44], and even generously assuming a total stock of 0.5 to 1.5 times GDP, aggregate land rent is 2 to 6 per cent of GDP. Complete confiscation of every krone of Danish land rent covers between 7 and 21 per cent of a 28.94 point gap. Reaching the gap from land alone at a 3 per cent tax on assessed value would require a pre-tax land stock of nearly seventeen times GDP.

Endowment. The requirement is 5.79 times GDP at a 5 per cent draw, 7.24 times at 4 per cent, and 11.58 times under the stewardship rule. For scale: Norway’s Government Pension Fund Global, the largest sovereign fund ever assembled, was worth NOK 21,268 billion at the end of 2025, about 4.84 times Norwegian mainland GDP, and took thirty years of North Sea petroleum revenue to build [40]. Denmark would need a fund between 1.2 and 2.4 times larger relative to its economy than Norway’s, with no petroleum. Corollary 1 gives accumulation times of 110 years at a 2 per cent saving rate and 77 years at 5 per cent.

Verdict. The programme fails for Denmark, and it fails by a wide margin rather than a narrow one. The consumption route requires 138 per cent tax-exclusive on the household base and 78 per cent on the broadest base its own VAT reaches; a rebate pushes the first figure to 173. Combining a perfectly enforced consumption tax at rates well above any observed, a land tax at the theoretical Georgist maximum and full carbon pricing still leaves eighteen to twenty-two percentage points of GDP unfunded. The honest conclusion is that a country which has chosen to socialise a very large share of consumption through the public budget cannot also decline to tax income, because it has spent the base it would need. Denmark’s household consumption is 43.6 per cent of GDP against America’s 67.9, and that single ratio does most of the work in the comparison. Denmark is not a counterexample that the framework can absorb with a footnote. It is a demonstration that Axiom 1 is conditional on the size and financing structure of the state.

5.3 Norway

Norway is a partial working proof, and it demonstrates both that the mechanism functions and how much it takes to make it function.

Norway collected 41.63 per cent of GDP in tax in 2023 and 40.19 in 2024. On the raw figures the productive-income aggregate is 30.88 per cent of GDP, which is 74 per cent of all Norwegian tax revenue and the highest dependence in our group.

That number is misleading, and correcting it is the whole Norwegian story. A large part of Norwegian corporate income tax is not a tax on production at all. Petroleum extraction faces a marginal rate of 78 per cent, combining the ordinary corporate rate with a special petroleum

Table 3: Norway: revenue composition, 2023 and 2024, as a percentage of total GDP including offshore petroleum [45]. Category 1000 is the published income-tax aggregate; because Norway’s category 1300 is zero it equals 1100 plus 1200 exactly, and it is the figure that must be used when summing to the total.

Base	2023	2024
Personal income tax (1100)	9.60	10.00
Corporate income tax (1200)	12.00	9.77
<i>income taxes combined (1000)</i>	<i>21.59</i>	<i>19.76</i>
Social security contributions (2000)	9.23	9.42
Payroll (3000)	0.05	0.05
Property (4000), of which recurrent 0.35	1.17	1.20
Goods and services (5000), of which VAT 7.40	9.58	9.75
Total	41.63	40.19
Productive-income aggregate	30.88	29.24

tax, and in economic substance this is a resource rent charge on a scarce natural endowment. That is pillar two of (2), which the reform keeps, rather than something it proposes to abolish.

The size of the reclassification can be recovered from three published ratios. Total tax is 41.63 per cent of total GDP; total tax over mainland GDP is 53.0 per cent; and total tax net of petroleum taxes over mainland GDP is 41.3. The first two give the mainland share of output, $41.63/53.0 = 78.5$ per cent. Substituting into the third, petroleum taxes are $41.63 - 41.3 \times 0.785 = 9.19$ per cent of total GDP. That is roughly three quarters of the entire 2023 corporate income tax line. Reclassifying it brings the effective Norwegian gap to 21.69 per cent of GDP, below Denmark’s and only three points above the United States’, in a country collecting sixteen points more tax.

The behaviour of the corporate line over time confirms the diagnosis rather than resting on the derivation. Norwegian corporate income tax was 18.28 per cent of GDP in the 2022 energy price spike, 12.00 in 2023 and 9.77 in 2024. A profits tax that halves in two years while the domestic economy grows is not taxing domestic profit; it is tracking a commodity price. The OECD attributes the 1.4 point fall in Norway’s total tax ratio between 2023 and 2024 to this line alone.

The fund. The Government Pension Fund Global stood at NOK 21,268 billion at the end of 2025, about 4.84 times mainland GDP [40, 41]. It fell to NOK 19,998 billion by the end of March 2026, or 4.55 times mainland GDP, which is worth stating because a fund of this size moves faster than the fiscal system it supports. The guideline targets a draw equal to the expected real return of 3 per cent, with actual withdrawals of 2.7 to 2.8 per cent in 2024 and 2025. At a 2.75 per cent draw on 4.84 times mainland GDP, transfers are 13.3 per cent of mainland GDP; at the full guideline they would be 14.5. The IMF records fund transfers as

covering roughly one fifth of central government budget expenditure [22].

The realised performance carries most of the empirical weight behind the endowment mechanism. The fund has earned an annual net real return of about 4.3 per cent since 1998, comfortably above the 3 per cent the fiscal rule assumes and above the draw actually taken. Norway has been more conservative than its own rule permits and has still seen the fund grow. Against that, forward-looking capital market assumptions from major asset managers currently imply only 2.7 to 3.2 per cent real for global equities, so the historical record should be read as the optimistic case rather than the central one.

The arithmetic, with the denominators kept straight. This is the point at which Norway’s two national accounts become dangerous, and our first pass through it made exactly the error the UAE section warns against. The reclassified gap of 21.69 per cent is measured against *total* GDP. The fund transfer of 13.30 per cent is measured against *mainland* GDP. Subtracting one from the other is not a quantity.

Converted to a common basis using the mainland share of 78.5 per cent derived above, the residual is

Basis	Residual
Both in mainland GDP	$27.61 - 13.30 = 14.31$
Both in total GDP	$21.69 - 10.45 = 11.24$

so the residual is eleven to fourteen per cent depending on the denominator chosen, not the eight to nine our earlier arithmetic produced. Adding petroleum rent proper at 9.19 per cent of total GDP would appear to close it, but counting both the rent and the fund double-counts the transition: the rent is precisely the revenue scheduled to disappear, and the fund exists to replace it rather than to supplement it.

A residual of eleven to fourteen points of GDP is not a gap that further accumulation bridges over a decade or two. It is the same order as the American gap. Norway is better placed than Denmark and it is not in a materially different position from the United States once the petroleum rent is set aside, which is a less comfortable conclusion than the one we drew before the denominators were reconciled.

The caveat. Norway did not accumulate this fund out of thrift applied to a normal economy. It accumulated it by discovering petroleum and declining to spend the proceeds, a combination available to very few countries and, at this point in the transition away from hydrocarbons, available to fewer still. Norway is a proof that the mechanism works. It is not a proof that the mechanism is generally reachable. The relevant question for every other country is whether

artificial intelligence can play the role that the North Sea played for Norway, and that question is addressed rather than answered in Section 7.

Verdict. Partial success. Norway could plausibly abolish personal income taxation over a multi-decade horizon while retaining consumption taxes and resource rents, and its own fiscal architecture already embodies most of the proposed design. It could not do so today, and it could not have done so without oil.

5.4 Singapore

Singapore is where the arithmetic nearly closes, and the reason it nearly closes is instructive about what the proposal actually costs.

Table 4: Singapore: revenue composition, FY2025 revised estimates [33]. Amounts are as published; GDP is implied by the reported total revenue ratio of 19.8 per cent, giving approximately S\$800 billion. Operating revenue excludes the Net Investment Returns Contribution; total revenue includes it.

Line	S\$ billion	% of GDP
Corporate income tax	35.24	4.41
Net Investment Returns Contribution	27.53	3.44
Goods and services tax (9 per cent)	21.30	2.66
Personal income tax	20.64	2.58
Operating revenue (excl. NIRC)	130.86	16.36
Total revenue (incl. NIRC)	158.39	19.80
Overall expenditure	143.29	17.91
Productive-income aggregate	55.88	6.99

Singapore’s productive-income gap is 6.99 per cent of GDP, well under half the American figure and less than a quarter of the Danish one. Two features explain the difference. The state is small relative to output, spending 17.9 per cent of GDP. And the Central Provident Fund is a mandated individual savings scheme rather than a tax, so Singapore carries none of the social contribution burden that adds six points of GDP in the United States and is folded into the income tax in Denmark.

The endowment is already there. The Net Investment Returns Contribution is the second-largest line in the Singapore budget at 3.44 per cent of GDP, funding roughly a fifth of government spending. The ranking is often reported wrongly: corporate income tax has exceeded NIRC every year since FY2022, and NIRC was the largest line only in FY2020 and FY2021 when corporate receipts collapsed. NIRC is 17.4 per cent of total revenue.

Under the Net Investment Returns framework, written into Article 142 of the Constitution,

the government may spend up to half of the expected long-term real returns on the net assets managed by GIC, by the Monetary Authority of Singapore, and, since April 2016, by Temasek [49]. Spending beyond that requires the concurrence of the elected President acting in his own discretion. The other half is reinvested, which is what preserves the reserves for future generations.

The arithmetic follows immediately, and it is startling. If 3.44 per cent of GDP is half the expected real return, the full expected real return is 6.88 per cent of GDP. The gap to be closed is 6.99. Singapore could therefore finance the abolition of *both* personal and corporate income tax almost entirely by spending the whole of its expected investment return rather than half of it, leaving a residual of 0.11 percentage points of GDP. At the observed yield of 0.296 points of GDP per point of GST, that residual is covered by raising the goods and services tax by about a third of a point, from 9 per cent to roughly 9.35. Singapore could stop taxing income altogether and the visible price would be a third of a point on the consumption tax.

What it costs. This is where the result turns, and the smallness of the residual makes the turn sharper rather than softer. Spending the full expected real return is precisely the move from the stewardship rule to the consumption rule described in Section 3. The fund's real value is preserved but its ratio to GDP decays at the growth rate. At 2 per cent real growth the endowment halves relative to the economy in about 35 years, and the Net Investment Returns Contribution falls with it, from 6.88 per cent of GDP back toward 3.44. The gap reopens within a generation, at which point the country faces the same choice with a smaller fund and no income tax administration left to restart quickly.

The 50 per cent rule exists for exactly this reason. It is Singapore's institutional expression of the principle that Paper V states as Article XII, intergenerational stewardship, and the constitutional requirement of presidential concurrence exists to stop precisely the manoeuvre that our arithmetic recommends. Applying Proposition 1 under the stewardship rule, permanent financing of a 6.99 point gap needs an endowment of 2.80 times GDP. The reserves implied by that expected real return at a 4 to 4.5 per cent portfolio return are 1.53 to 1.72 times GDP, which is 55 to 61 per cent of the stewardship requirement.

The gap between those two figures is the entire content of the Singaporean result. On a consumption rule Singapore can abolish income taxation today. On a stewardship rule it is a little over half way there. The 0.10 point residual is not a measure of how close the proposal is to working; it is a measure of how much a country can borrow from its descendants without the transaction appearing anywhere in the current budget.

Verdict. Singapore is the only country in our set where the model closes on current resources, and it closes only by converting a permanent endowment into a several-decade annuity. A

framework whose own constitutional text forbids externalising liabilities onto future generations cannot coherently recommend this. The defensible Singaporean version of the proposal is slower: abolish the personal income tax now, financed from a modest GST increase and the existing NIRC, and continue accumulating toward the 2.80 multiple before touching the corporate tax. Corollary 1 applied from Singapore’s current position rather than from zero puts the remaining accumulation at about twelve years if the state contributes 5 per cent of GDP annually, or sixteen at 2 per cent, starting from reserves of roughly 1.6 times GDP. That is the shortest transition path of any country studied here by a wide margin.

5.5 United Arab Emirates

The UAE is the closest realised analogue to $\tau_{\text{labor}} = 0$ that exists. There is no personal income tax. A 9 per cent federal corporate tax was introduced only in June 2023, with a 15 per cent rate for large multinationals from January 2025, and a 5 per cent VAT dates from 2018. General government revenue is nonetheless 28.5 per cent of GDP in 2023 and 27.8 in 2024, comparable to the United States, and the country ran a fiscal surplus of 6.4 per cent of GDP in 2024 [23].

It works. The question is what it is resting on.

Hydrocarbon revenue is 47.0 per cent of government revenue in 2023 and 53.1 in 2024, which is 13.4 and 14.8 per cent of GDP respectively. VAT and excise together raised about AED 41 billion in 2024 and over 46 billion in 2025; the two are published only as a combined figure. We report no yield for the 9 per cent corporate tax because no primary figure exists: the IMF books zero corporate income tax in its 2023 to 2025 outturns, and the figure of about 3.3 per cent of GDP that circulates in secondary commentary is a 2021 estimate for a hypothetical 10 per cent tax rather than a measurement of the enacted one.

The decisive figure, and the one that should be quoted whenever the UAE is offered as proof that income taxes are unnecessary, is the adjusted non-hydrocarbon primary deficit: 17.1 per cent of *non-hydrocarbon* GDP in 2024, which is about 13.3 per cent of total GDP. The denominator matters and is often dropped. Strip out the oil and the UAE is not a balanced low-tax state. It is a state running a very large structural deficit that a resource rent happens to cover.

The endowment test. This is the case where the proposed architecture has the most to say, because converting a depleting rent into a permanent endowment is exactly what (5) describes. Replacing 13.4 per cent of GDP of hydrocarbon revenue permanently requires an endowment of 2.68 times GDP at a 5 per cent draw, 3.35 times at 4 per cent, and 5.36 times under the stewardship rule.

The asset side cannot be assessed with the precision the question deserves. Mubadala does disclose, at \$330 billion at the end of 2024 and \$385 billion at the end of 2025. ADIA does not

disclose at all; the most commonly cited external estimate is about \$1.19 trillion, and published estimates span roughly \$0.9 to \$1.2 trillion. ADQ is estimated at about \$251 billion on the authority of a single spokesperson statement. Taken at face value, and the caveat is substantial enough that we would not defend any conclusion resting on it, those estimates would put the UAE somewhere near the endowment a permanent replacement of its hydrocarbon revenue would require. We state this as a possibility that the public data cannot settle rather than as a finding. It could be settled quickly by disclosure, and the fact that it cannot be settled without disclosure is itself relevant to the governance objection in Section 8.

Verdict. The UAE proves less than it appears to and more than it is usually given credit for. It does not prove that a state can function without taxing production; it proves that a sufficiently large rent can substitute for taxing production while the rent lasts. But it is also the clearest existing demonstration of the transition the architecture describes, converting an exhaustible resource rent into a diversified financial endowment at roughly the scale the arithmetic requires. Whether the conversion has actually been completed cannot be verified from public data, and any confident claim in either direction is unsupported.

5.6 Cross-country summary

6 The Consumption Replacement

6.1 Rates, and the two ways of quoting them

A consumption tax can be quoted two ways and the choice changes the number a reader hears without changing the tax. Let τ_e be the rate applied to the pre-tax price, the convention used for every value added tax in the world. Let τ_i be the same tax expressed as a share of the final price paid. Then

$$\tau_i = \frac{\tau_e}{1 + \tau_e}, \quad \tau_e = \frac{\tau_i}{1 - \tau_i}. \quad (11)$$

A 25 per cent Danish VAT takes 20 per cent of the gross price. A tax advertised at 23 per cent tax-inclusive is 29.9 per cent tax-exclusive. The two conventions diverge without limit as rates rise, which is why a debate conducted in mixed units cannot be settled. We give both throughout.

To raise revenue R as a share of GDP from a base B , also as a share of GDP, $\tau_e = R/B$. The base is household final consumption net of the consumption taxes already in it, multiplied by a coverage share κ capturing exemptions, reduced rates and evasion. Two assumptions in this construction pull the reported rates in opposite directions, and since this paper argues for the policy whose price it is computing we state the net rather than only the flattering half.

Subtracting the whole of existing goods-and-services tax revenue from the base is conservative:

Table 5: Replacement gaps and required public capital endowments. Endowment ratios follow Proposition 1; the stewardship column uses $\rho = r - g = 2.5\%$. Accumulation times follow Corollary 1 at a saving rate of 5 per cent of GDP, starting from zero, and therefore overstate the remaining task for Singapore, Norway and the UAE, which have already accumulated.

Country	Gap (% GDP)	Required W_G/Y			Years to accumulate
		$\rho = 5\%$	$\rho = 4\%$	$\rho = 2.5\%$	
United States	18.51	3.70	4.63	7.40	62
<i>excl. social contributions^a</i>	12.48	2.50	3.12	4.99	50
Denmark	28.94	5.79	7.24	11.58	77
Norway, as reported ^b	30.88	6.18	7.72	12.35	79
<i>after petroleum reclassification</i>	21.69	4.34	5.42	8.68	67
Singapore ^c	6.99	1.40	1.75	2.80	35
UAE ^d	13.39	2.68	3.35	5.36	52

All gaps are percentages of GDP, not of total tax revenue. OECD countries are 2023, Singapore FY2025, the UAE 2023.

^a Robustness row treating social security contributions as contributory deferred compensation rather than as a tax on production, and therefore leaving them in place. See Section 4.

^b The second Norwegian row removes petroleum taxation, derived at 9.19 per cent of GDP in Section 5.3, and counts it as resource rent, which is its economic substance. Norway has also already accumulated a fund of 4.84 times mainland GDP, so its remaining task is far smaller than the accumulation column implies.

^c Singapore already holds reserves of roughly 1.6 times GDP, so its remaining accumulation is about twelve years rather than 35.

^d Hydrocarbon revenue rather than income tax, since the UAE has no personal income tax. This is the rent that must be made permanent.

Table 6: What each pillar can carry, in percentage points of GDP, against a best observed performance anywhere. Theoretical ceilings assume complete capture with no exemptions and no behavioural response; they are upper bounds, not forecasts, and are not additive without double counting.

Pillar	Best ever observed	Observed	Theoretical
Land value	Denmark <i>grundskyld</i> 1.09 (2023), falling to 0.77 (2024); Estonia’s pure land tax 0.15	1.1	4–11 ^a
All recurring property	UK 2.80 (2024); UK 4.18 (1982) is the all-time OECD maximum	2.8	4.2
Resource rents	UAE 14.8 (2024); Norway petroleum 9.19 (2023)	14.8	requires a resource
Externalities ^b	Norway carbon tax 0.44; no country has ever exceeded 1	0.44	3.4 (US)
Final consumption	Hungary’s 27 per cent rate; Denmark raises 9.14	9–10	base-limited
Public capital	Norway 13.3 of mainland GDP	13.3	$\rho \cdot W_G/Y$

^a Full capture of aggregate land rent $R_L = iV_L$. The range spans the United States at the bottom, with land of 1.0–1.6 times GDP, and the higher advanced-economy land stocks of 2.4–2.8 times GDP. See Section 7.2.

^b Transitional, not structural. A Pigouvian charge set at the social cost is designed to shrink its own base, so this row should be read as a declining contribution over the second half of the century rather than as permanent revenue. See Section 7.8. We include it at its theoretical maximum in the country tests, which is generous, and it changes no verdict.

some excise falls on business inputs and government purchases rather than on household consumption, so the true household base is somewhat larger than we assume and the reported rates are correspondingly too high. Working the other way, and more powerfully, $\kappa = 0.96$ is borrowed from a system whose revenue ratio is computed on a base that includes government consumption. A tax confined to household consumption has fewer easily-reached transactions and a worse coverage ratio than one that also invoices the state, so 0.96 on our narrower base is optimistic.

The second effect is the larger of the two. Moving κ from 0.96 to 0.90 raises the American requirement from 37.2 to 39.7 per cent, whereas a two-point widening of the household base lowers it by about 1.1 points. The net bias in the headline figures is therefore toward understatement: the true required rates are more likely above those in Table 7 than below them. Readers should treat that table as a floor. The empirical anchors for κ are the OECD's VAT Revenue Ratios for reference year 2022: New Zealand reaches 0.96 on a genuinely single-rate base, the OECD average is 0.58, and Denmark, despite a 25 per cent headline, reaches 0.66 (43, §2.9 and Annex Table 2.A.7). The gap between 0.96 and 0.58 is the cost of exemptions and reduced rates, and it is worth roughly two thirds of a tax system.

One denominator caution, because getting it wrong reproduces neither our figures nor the OECD's. The VAT Revenue Ratio is computed against total final consumption, which *includes government consumption*, whereas κ here is defined against household final consumption alone. The two are not the same quantity and a ratio measured on one cannot be transferred to the other. We use 0.96 as an indicative ceiling on what a clean single-rate design can achieve rather than as a measurement carried across bases, and part of the reason New Zealand's ratio is so high is precisely that its GST reaches government purchases which our household base excludes. Read $\kappa = 0.96$ as coverage as good as the best-designed system in the world, not as New Zealand's number applied here. The OECD has published no VAT Revenue Ratio for any country after 2022.

6.2 What each country would have to charge

Table 7 is the central result of the paper and it divides the group cleanly in two.

The United States and Singapore land at 37.2 and 35.3 per cent tax-exclusive on a clean base, which is 27.1 and 26.1 per cent tax-inclusive. These are large rates. They are not absurd ones. They sit about ten points of tax-exclusive rate above Hungary's 27 per cent, the highest standard rate in force anywhere, and they would apply to a base far broader than Hungary's. A country willing to adopt New Zealand's base discipline and to accept a rate no European state has tried could abolish every tax on earning, immediately, without touching land or carbon or waiting on a sovereign fund.

Denmark and Norway cannot. Denmark needs 138 per cent tax-exclusive on the household

Table 7: Consumption tax required to abolish all taxes on labour income, corporate profit and social insurance, and to replace existing consumption taxes at the same time. R is the total revenue the consumption tax must raise. The household base is final consumption net of existing consumption taxes. $\kappa = 0.96$ is New Zealand’s coverage; $\kappa = 0.58$ is the OECD average. Rates are given tax-exclusive with the tax-inclusive equivalent in brackets.

Country	R	Household base (% GDP)	Required rate, excl. (incl.)	
	(% GDP)		$\kappa = 0.96$	$\kappa = 0.58$
United States	22.76	63.65	37.2% (27.1%)	61.7% (38.1%)
Singapore ^c	9.65	28.44	35.3% (26.1%)	58.5% (36.9%)
Norway, reclass. ^{a,b}	31.27	32.92	98.9% (49.7%)	163.8% (62.1%)
Denmark	41.35	31.19	138.1% (58.0%)	228.6% (69.6%)

^a After reclassifying petroleum taxation as resource rent, per Section 5.3. On the unadjusted Norwegian gap the $\kappa = 0.96$ rate is 128.0 per cent tax-exclusive.

^b The Norwegian household base rests on household final consumption of 42.50 per cent of total GDP, which is the one consumption input in this table we could not verify against a primary series; it is an implied figure and it should be treated as such. The result is sensitive to it. At 40 per cent the required rate is 107.1 per cent tax-exclusive and at 36 per cent it is 123.3, against the 98.9 shown. A reader who prefers a lower Norwegian consumption share should read the Norwegian verdict as correspondingly worse, never better.

^c $\kappa = 0.58$ is not meaningful for Singapore. Its observed yield of 0.296 points of GDP per rate point implies an effective base of 29.56 per cent against a net household base of 28.44, that is an implied coverage above one, because the GST reaches tourist and business consumption that the household base excludes. The 58.5 per cent figure is reported for symmetry only and should not be read as a Singaporean scenario.

base. Widening the base helps less than one would hope. Denmark’s value added tax is assessed against a theoretical base of 55.4 per cent of GDP, derived from its published revenue ratio, and that base includes government consumption. At New Zealand’s coverage of that wider base the requirement is 77.8 per cent tax-exclusive and 43.7 per cent tax-inclusive; at Denmark’s own current coverage of 0.66 it is 113.1 per cent. Two cautions apply to the wider base and both cut against it. Government consumption cannot be taxed without raising the government’s own costs, so a requirement computed on a base that includes it is understated unless the revenue target is grossed up correspondingly, which we have not done. And the theoretical base already carries Denmark’s exemption structure in its denominator, so applying a further coverage factor to it is a statement about a reformed Danish VAT rather than about the one that exists. Norway after petroleum reclassification needs 98.9 per cent on the household base and 66.0 per cent on the broad one. Rates in that region do not raise the revenue attributed to them, for reasons Section 6.4 takes up.

The asymmetry has one dominant cause and it is not the tax system. American households consume 67.9 per cent of GDP; Danish households consume 43.6. Denmark routes an unusually large share of lifetime consumption through the public budget rather than through households, which shrinks exactly the base a consumption tax needs while enlarging exactly the expenditure it must fund. A country cannot socialise consumption and then tax consumption to pay for it.

6.3 Reconciling with a coarser calculation

Figures previously circulated from this exercise gave a different pair of numbers for the same question, and since they are in circulation we reconcile them rather than quietly replacing them. Those figures put the United States at 61 per cent on an observed-coverage base and 34 per cent on a perfectly covered one, and Denmark at 104 and 87. *Three* changes account for the difference, and it matters that there are three, because on the two we originally disclosed the reconciliation does not work.

The first is the base. The circulated calculation applied the rate to household final consumption at purchaser prices, which already contains the consumption taxes being replaced. A tax cannot be levied on itself, so the base must be net of them. Netting removes 4.25 points of GDP in the United States and 12.41 in Denmark.

The second is what the new tax replaces. The circulated calculation replaced the income taxes plus the value added tax; this one replaces the income taxes plus the whole of goods and services taxation, which is the honest reading of a single consolidated consumption tax. For Denmark that raises the requirement from 38.08 to 41.35 points of GDP. For the United States the change is null by construction, since there is no federal VAT and the circulated calculation already used the whole 4.25 line.

The third is the coverage parameter, and it is the one that does the work. The circulated figures used $\kappa = 0.55$ for the observed case and $\kappa = 1.00$ for the covered one; this paper uses the OECD average of 0.58 and New Zealand's 0.96.

The consequence is that our earlier characterisation of the American reconciliation was wrong, and we correct it here rather than leaving it standing. Applying only the first two changes moves the United States from 61.0 to 65.0 per cent. The return to 61.7 comes entirely from the change in κ . It is therefore not true, as we previously wrote, that the American figures are barely affected and that 61 and 61.7 are the same number arrived at twice; the agreement is a coincidence of two offsetting revisions.

A like-for-like comparison, holding coverage at full in both vintages, is the honest test of robustness, and it passes: the United States moves from 33.5 to 35.8 per cent, and Denmark from 87.3 to 132.6. The American result is genuinely insensitive to the base corrections. The Danish result is not, and the circulated figure understated Denmark by about forty-five points of tax-exclusive rate.

6.4 The rate at which a base stops behaving

Every figure in Table 7 assumes the base survives the rate, and above some level it does not. Murray's analysis of a national retail sales tax concluded that evasion and avoidance become first-order at rates far below those in the Danish column, because a retail tax collects the entire

liability at a single point from a single agent with an obvious incentive and no cross-check [38]. A value added tax is more robust, since the invoice-credit chain makes each trader’s claim a report on his supplier’s sale, but the protection weakens at the final stage, which is where the whole of the tax rests.

Gale’s calculation for the United States is the closest published analogue to Table 7 and it points the same way [14]. Working on a narrower replacement than ours, federal taxes only, he found required rates well above those advertised by proponents once realistic allowances were made for evasion, for government and non-profit consumption, and for the transition. Our figures are not directly comparable to his, since we replace taxes at all levels of government and include existing consumption taxes in the total, which raises the requirement.

The best-known American proposal of this kind states its rate in the statute rather than in commentary. The FairTax Act provides that “the rate of tax is 23 percent of the gross payments for the taxable property or service”, and defines gross payments to include the federal tax itself [58]. That is a tax-inclusive rate, and by (11) it is 29.9 per cent tax-exclusive. The conversion is the useful part of the comparison, because 23 and 30 describe the same tax and a debate conducted in mixed units cannot be settled.

We draw no conclusion from setting that 23 per cent against our own 27.1, and we previously did. The scopes do not match. H.R. 25 replaces federal income, payroll, estate and gift taxes; our requirement of 22.76 per cent of GDP replaces income and payroll taxes at *all* levels of government and additionally absorbs every existing state and local sales tax. Comparing the two rates is comparing a federal reform to a federal, state and local one, and the difference in scope runs in the direction that flatters our figure. A like-for-like comparison would need a decomposition of American income taxation into federal and sub-federal components, which we have not done and which our sources do not give us at the vintage we use elsewhere. Until that is done the honest position is that our arithmetic neither supports nor refutes the advertised rate, and we withdraw the claim that it sits below what the arithmetic requires.

What can be said without the decomposition is Gale’s, not ours: working on a federal-only replacement, and therefore on matched scope, he found required rates well above those advertised by proponents once evasion, government and non-profit consumption, and the transition were allowed for [14].

We therefore report Table 7 as an arithmetic lower bound on the required statutory rate rather than as a forecast of a workable one. The behavioural response is not in it. For the United States and Singapore that response is a question about the design of enforcement. For Denmark and Norway it is not, because no enforcement regime recovers a base at a 138 per cent rate on an open border with Germany.

6.5 Regressivity, and the instrument that answers it

A consumption tax takes a larger share of a poor household’s annual income than of a rich one’s. This is the standard objection and it is correct as stated. Two things answer it, one analytical and one a design choice with a price attached.

The analytical point is the budget-constraint identity set out in Section 2, and it should be attributed there rather than to Atkinson–Stiglitz, whose theorem concerns differential rates across goods and not the choice between taxing consumption and taxing income. Over a lifetime, and absent bequests, a uniform consumption tax and a flat labour income tax fall on the same base, because lifetime labour income is either consumed or bequeathed. Annual regressivity and lifetime proportionality are both true and they are measuring different things. What a uniform consumption tax genuinely gives up relative to a graduated income tax is the ability to vary marginal rates with ability. The average burden survives; the graduation does not. The honest response is to recover that progressivity somewhere else and to say plainly that it was lost.

The instrument is a universal rebate paid to every household, sized to refund the tax on a subsistence consumption bundle. It does not vary with income and so creates no additional marginal rate, which is what distinguishes it from a means-tested transfer. Its cost is that the statutory rate must rise to fund it. Table 8 prices it.

Table 8: Statutory rate required when a universal rebate refunds the tax on the first s of average consumption, at $\kappa = 0.96$. Tax-inclusive equivalents in brackets.

Country	$s = 0$	$s = 15\%$	$s = 20\%$
United States	37.2% (27.1%)	43.8% (30.5%)	46.6% (31.8%)
Singapore	35.3% (26.1%)	41.6% (29.4%)	44.2% (30.6%)
Denmark	138.1% (58.0%)	162.5% (61.9%)	172.6% (63.3%)

The rebate converts a proportional tax into a progressive one on consumption. A household consuming X times the subsistence bundle pays an effective tax-inclusive rate of $\tau_i(1 - 1/X)$, which is negative below the threshold, zero at it, and rises toward τ_i thereafter. At the American rate with a 20 per cent rebate the schedule runs from -31.8 per cent at half the subsistence bundle, through zero at subsistence and 15.9 per cent at twice it, to 25.4 per cent at five times and 28.6 per cent at ten.

Two honest qualifications. That schedule is progressive in consumption, not in income, and the two diverge exactly where the very rich save most of what they receive; a consumption tax reaches a fortune only when it is spent, and some fortunes are never spent by the person who accumulates them. And the rebate raises the American required rate by 9.3 points of tax-exclusive rate. That is the price of the distributive property, and it is large enough that it has to be quoted alongside it.

7 What the Arithmetic Shows

7.1 The rent ceiling

The most robust empirical finding in this paper is a ceiling rather than a possibility. Across every country and every pillar, the bases that the proposal wants to tax are smaller than the base it wants to stop taxing, and they are smaller by a factor of roughly three to five in the hard cases.

This is not an accident of the current tax system. It reflects the structure of national income. Land rent, resource rent and measurable externality damage together account for a modest share of output in advanced economies, while labour compensation accounts for somewhere around 55 to 65 per cent of it. A tax architecture that renounces the largest component of national income and taxes only the smallest components must either accept a much smaller state or find an entirely different kind of revenue. The proposal's answer is the endowment, which is why the endowment does all the work and why the paper spends most of its space on it.

7.2 What the Henry George Theorem actually delivers

The preceding subsection understates the Georgist case, and correcting the understatement produces the sharpest result in the paper.

Arnott and Stiglitz proved that in a city of optimal size financing a pure local public good, aggregate land rents exactly equal expenditure on that public good [3]. The result, generally called the Henry George Theorem, is not a slogan but a theorem, and it says something quite specific: the benefit of a non-rival, spatially fixed amenity is capitalised into the value of the land that enjoys access to it, so the rent it generates is precisely sufficient to pay for it.

Set that against the observed record and there is an apparent contradiction. The highest verified pure land tax yield anywhere is Denmark's *grundskyld* at 1.09 per cent of GDP in 2023, falling to 0.77 in 2024. Estonia, which operates the purest land value tax in the world with no building tax and no property transfer tax at all, raises 0.15 per cent of GDP from it. Meanwhile public spending in these countries runs from eighteen to forty-five per cent of GDP. Either the theorem is wrong or something has been measured incorrectly.

Something has been measured incorrectly, and the details matter. Observed yields are low because observed land taxes are levied at low rates on assessed *value*, not because land rent is small. If land is worth V_L before tax and the market yield on it is i , the annual rent is $R_L = iV_L$. Advanced-economy land stocks run from about 1.0–1.6 times GDP in the United States to 2.4–2.8 in France, the United Kingdom and Japan, and 3.8 in Australia [31, 44]. At a four per cent yield that is an aggregate land rent of roughly 4 per cent of GDP at the American

end of the range, 11 at the French, British and Japanese end, and 15 in Australia. Denmark’s *grundskyld* at one per cent of GDP is not capturing one per cent of the available rent; on a land stock of one times GDP it is capturing about a quarter of it, and on a larger stock, less. The Georgist ceiling is several times the observed yield, and our earlier framing was unfair to George.

A subtlety here is easy to get wrong and we state it explicitly because it changes the numbers. A tax levied at rate τ_L on *assessed market value* does not collect $\tau_L V_L$, because the tax itself depresses the value against which it is assessed. From the capitalisation result in Appendix A.3, the post-tax value is $V' = R_L/(i + \tau_L)$, so revenue is

$$\tau_L V' = R_L \cdot \frac{\tau_L}{i + \tau_L}, \quad (12)$$

and the share of rent captured is $\tau_L/(i + \tau_L)$. Setting $\tau_L = i$, which sounds like full capture, collects exactly half the rent. Capturing 75 per cent requires $\tau_L = 3i$, and complete capture is approached only asymptotically as the value base collapses toward zero. Table 9 gives the schedule.

Table 9: Share of land rent captured by a tax on assessed market value, at a four per cent market yield. Revenue is $R_L \tau_L/(i + \tau_L)$; the third column illustrates the case of aggregate land rent of 8 per cent of GDP, corresponding to a pre-tax land stock of twice GDP. That is near the middle of the advanced-economy range and above the American level, so the revenue column should be scaled down by roughly a third for the United States.

Tax rate τ_L	Share of rent captured	Revenue (% of GDP)
1%	20.0%	1.60
2%	33.3%	2.67
3%	42.9%	3.43
4%	50.0%	4.00
8%	66.7%	5.33
12%	75.0%	6.00
Assessment on annual rental value instead		8.00

This is a practical argument for assessing land on its annual rental value rather than its capital value, which is what the United Kingdom does for business rates and what a serious Georgist programme would have to do. A rental-value assessment at 100 per cent of rent collects R_L directly and does not suffer from a base that shrinks as the rate rises. The ceiling is reachable in principle, but only with the right instrument, and none of the countries in this study uses it.

Denmark’s recent experience adds a second practical warning. Its land tax yield fell by a quarter between 2023 and 2024, not because land became less valuable but because a statutory reform moved assessments onto a market basis and cut municipal rates. The Georgist claim that land cannot flee taxation is a claim about the asset. It is not a claim about the legislature,

and the legislature is the part that has historically moved.

That correction makes the failure diagnosis much more precise, though it is kinder to George in the countries with expensive land than in the countries without. Full rent capture yields roughly 10 to 11 per cent of GDP where land is worth 2.4 to 2.8 times output, which is close to what advanced states actually spend on public goods in the economist's sense: defence, courts, policing, general administration, basic research, and the non-transfer part of infrastructure and education. That category runs to about twelve to fourteen per cent of GDP in the countries studied here. In those economies the Henry George Theorem is roughly right and a full Georgist tax could finance most of the public-goods state. In the United States, where land is only 1.0 to 1.6 times GDP, full capture yields 4 to 6.4 per cent and covers perhaps a third to a half of it.

What it cannot finance is the transfer state. Social insurance, pensions, health insurance and income support account for the majority of expenditure in every country in this study, and they do not capitalise into land values in the way the theorem requires. A pension is not a spatially fixed amenity. Its value does not attach to a location, so no land rent arises to pay for it. Denmark fails hardest in our tests for exactly this reason and not because Danish land is unusually cheap: Denmark has chosen to route an unusually large share of lifetime consumption through the public budget, and that share is transfers.

The finding can therefore be stated with more force and more precision than the crude version. The proposed architecture is a complete financing system for a state that provides public goods. It is not a financing system for a state that redistributes income. Whether to have the second kind of state is a political question that this paper takes no position on, but the arithmetic settles the conditional: a country that wants both a large transfer state and no tax on production must own the capital that pays for the transfers, which returns the argument to the endowment and to Proposition 1.

7.3 The stock–flow confusion

Stated as a set of tax rates, Axiom 1 sounds like a policy that a legislature could adopt in a single session. Stated correctly, it is a claim about the composition of national wealth. Proposition 1 makes the translation: at a sustainable 2.5 per cent draw, abolishing one percentage point of GDP of production taxation requires the state to own 40 per cent of GDP in financial assets.

Piketty and Zucman document private wealth-to-income ratios in rich countries recovering to somewhere between four and seven times national income over the twentieth century [47]. Set that against the requirements in Table 5. The American stewardship requirement of 7.40 times GDP and the Danish requirement of 11.58 exceed the entire private wealth stock of those countries. The proposal, taken to its pure form and financed entirely from the endowment, is a proposal for the state to become the majority owner of national wealth.

That is a serious thing to have discovered, and it should change how the proposal is argued. A framework that positions itself as neither capitalist nor socialist, and which in Paper V protects private property against arbitrary confiscation and protects economic pluralism against capture, cannot casually contain a fiscal mechanism whose pure form requires the state to hold more assets than the private sector does. Either the pure form is abandoned in favour of a partial one, which is what Article VII's schedule and its emergency departure clause between them already allow, or the tension has to be faced directly.

7.4 The endowment lowers the rate toward a floor

An earlier framing of this argument treated the public capital endowment as a precondition: the state accumulates for decades, and only when the fund is large enough does abolition become possible. That framing is wrong in one direction and, as it turns out, too generous in another. Abolition does not wait on the endowment; but neither does the endowment ever finish the job.

Take the first point first. Section 6 shows the consumption switch is available immediately, at 37.2 per cent tax-exclusive for the United States. What a fund does is reduce that rate. Table 10 states the comparative static: given a fund of size W_G/Y already in hand, drawn at ρ , what rate must consumption carry?

Table 10: United States, comparative static: the consumption tax rate required if the state *already held* a fund of the stated size, drawn at 4 per cent, on a clean base ($\kappa = 0.96$, $B = 61.10$ per cent of GDP) with no rebate. This is not a path; reachability is the subject of Table 11. Note that a 4 per cent draw is neither of the two rules of Section 3: with $r = 4.5$ and $g = 2$ it exceeds $r - g$ by 1.5 points, so a fund held at this draw shrinks by 1.5 per cent a year relative to GDP. Every rate below is therefore a transient, not a stationary state.

Endowment W_G/Y	Fund revenue (% GDP)	Residual R	Consumption rate
none	0.00	22.76	37.2% (27.1%)
1×	4.00	18.76	30.7% (23.5%)
2×	8.00	14.76	24.2% (19.5%)
3.70×	14.80	7.96	13.0% (11.5%)
5.69×	22.76	0.00	0.0%

7.5 What the fund can actually reach

Table 10 answers a question about stocks. The question that matters is whether those stocks are reachable, and here our first attempt at this section was simply wrong in a way worth setting out, because the error is easy to make and it flattered the argument.

A fund that is being drawn on cannot be propagated by the accumulation law of Corollary 1. That corollary is derived with no withdrawals. If the state contributes a share s of GDP and

simultaneously withdraws a share ϕ of the fund, the correct law nets the two:

$$\dot{w} = s + (r - g - \phi) w. \quad (13)$$

Writing $\delta = \phi - (r - g)$ for the excess of the draw over the growth-adjusted return, the solution from $w(0) = 0$ is

$$w(t) = \frac{s}{\delta} \left(1 - e^{-\delta t}\right), \quad w(\infty) = \frac{s}{\delta}, \quad (14)$$

which for $\delta > 0$ is bounded. The fund does not grow without limit; it converges.

At the parameters used throughout this paper, $s = 5$ per cent of GDP, $r - g = 2.5$ per cent and $\phi = 4$ per cent, we have $\delta = 1.5$ per cent and the fund converges to $w(\infty) = 3.33$ times GDP. Maximum fund revenue is $\phi w(\infty) = 13.33$ per cent of GDP. Since the consumption tax must raise $22.76 + 5.00 = 27.76$ per cent of GDP while the contribution continues, the residual cannot fall below 14.43 per cent of GDP, and the rate cannot fall below

$$\tau_e^{\text{floor}} = \frac{14.43}{61.10} = 23.6 \text{ per cent} \quad (19.1 \text{ per cent tax-inclusive}). \quad (15)$$

Table 11: United States: the consumption tax rate along a transition that contributes 5 per cent of GDP a year and draws 4 per cent of the fund a year, propagated by (13). The fund converges to 3.33 times GDP and the rate converges to 23.6 per cent tax-exclusive. Zero is not attained at any date.

Year	W_G/Y	Fund revenue	Consumption must raise	Rate, excl. (incl.)
0	0.00	0.00	27.76	45.4% (31.2%)
10	0.46	1.86	25.90	42.4% (29.8%)
20	0.86	3.46	24.30	39.8% (28.5%)
30	1.21	4.83	22.93	37.5% (27.3%)
40	1.50	6.02	21.74	35.6% (26.2%)
60	1.98	7.91	19.85	32.5% (24.5%)
120	2.78	11.13	16.63	27.2% (21.4%)
∞	3.33	13.33	14.43	23.6% (19.1%)

Three things follow, and none of them is what we expected.

The transition is far slower than the comparative static suggests. The rate does not fall below the no-contribution figure of 37.2 per cent until year thirty-one, not year nineteen. After a century it is still above 28 per cent.

The rate never reaches zero at these parameters. A zero rate needs fund revenue of 27.76 per cent of GDP, which at a 4 per cent draw requires a fund of 6.94 times GDP; the fund converges to 3.33. No amount of patience closes that gap, because the shortfall is structural rather than temporal.

Lowering the draw does not rescue the result on any horizon a legislature could act on. Reducing ϕ raises the asymptote, since $w(\infty) = s/(\phi - (r - g))$, but it also slows convergence,

since δ falls in the same proportion. At $\phi = 3$ per cent the asymptote rises to 10 times GDP and zero does become attainable in principle, but the time constant is 200 years and the zero arrives after roughly five centuries. At the stewardship draw $\phi = r - g$ the fund grows linearly and reaches zero at about year 222. These are not policy horizons.

The honest statement is therefore weaker than the one we set out to make, and weaker than the version that appeared in earlier drafts of this argument. A public capital endowment lowers the consumption tax rate toward a floor. It does not abolish the consumption tax, and at plausible saving rates and returns it does not come close. The floor is worth having: 23.6 per cent tax-exclusive is a normal European rate on a broad base, against 37.2 without the fund. But the destination is a lower rate, not no rate, and any version of this proposal that promises the elimination of consumption taxation through public investment is promising something the arithmetic does not deliver.

7.6 Where the contribution comes from

The endowment still has to be funded, and the candidates matter because most of them violate the axiom. Taxing agentic profits to build the fund is an income tax on production. The defensible sources are the ones the architecture already contains: resource rents, spectrum and concession auctions, land tax revenue in excess of current spending, and public equity taken in exchange for public investment in research, compute or energy infrastructure. To these the consumption tax adds one more, which is simply a rate set above current requirements, with the excess directed to the fund.

Norway remains the instructive case and the caution. It accumulated without taxing its own citizens because the contribution came from a resource rent that was, in the relevant sense, free. A country without such a rent must fund the fund out of the consumption base. Whether artificial intelligence provides an equivalent windfall depends on whether the returns to autonomous productive capital accrue publicly or privately, and Section 2.3 has already argued that under free replication they accrue to inelastic inputs rather than to anyone in particular.

7.7 Growth makes it harder

An intuition worth dispelling: the AI productivity boom that motivates this series does not help, and the sign is unambiguous rather than a matter of parameters.

Under the stewardship rule the sustainable draw is $r - g$, so faster growth lowers the draw rate and raises the required endowment ratio $w^* = \gamma/(r - g)$. One might expect an offsetting effect, on the grounds that faster growth raises the absolute contribution available from a given saving share. It does, but the offset is exactly cancelled.

Proposition 2 (Growth lengthens accumulation). *Let the revenue share to be replaced, γ ,*

and the saving rate, s , both be constant shares of GDP, and let the target be the stewardship endowment $w^* = \gamma/(r - g)$. Then the accumulation time is

$$t^* = \frac{\ln(1 + \gamma/s)}{r - g}, \quad (16)$$

which is strictly increasing in g for $g < r$.

Proof. Substitute $w^* = \gamma/(r - g)$ into (9). The argument of the logarithm becomes $1 + w^*(r - g)/s = 1 + \gamma/s$, which does not involve r or g . The prefactor $1/(r - g)$ is strictly increasing in g on $g < r$. $\square$

The cancellation has an interpretation as well as a derivation. Faster growth raises the contribution flow and the target in the same proportion, because both are denominated in GDP; what it does not raise is the rate at which the fund outgrows the economy, and that is the only thing that compounds.

The result holds r fixed, and that assumption does real work. In a standard growth model the real return co-moves with the growth rate, so that faster growth raises r as well as g and leaves $r - g$ closer to constant than Proposition 2 assumes. If $r - g$ is invariant to g , the effect disappears entirely. We state the result because the opposite intuition, that an AI growth acceleration makes the endowment easy, is more common and is wrong under any parameterisation we can construct; but the honest form of the claim is that faster growth does not help, rather than that it necessarily hurts by the full amount in the table. The effect is large. For the United States at a 5 per cent saving rate, accumulation takes 44 years at 1 per cent growth, 62 years at 2 per cent, and 103 years at 3 per cent. A country that succeeds in raising its growth rate by a percentage point through automation lengthens its path to a production-tax-free fiscal system by four decades.

The result cuts directly against the framing in which AI abundance simply makes public finance easy.

7.8 The shrinking externality base

A structural point about pillar three that applies regardless of country. A Pigouvian tax set at the social cost of the externality is designed to reduce the externality. If it works, the base disappears. Carbon revenue at full social cost is worth about 3.3 per cent of American GDP today and is intended to approach zero over the second half of the century.

Building a permanent revenue architecture on a base whose deliberate purpose is self-elimination is a design error. Externality charges belong in the system for allocative reasons, and they should be levied whether or not revenue is needed, but they cannot be counted as

structural revenue in a long-run fiscal plan. We have included them in the country tests at their maximum theoretical yield, which is generous, and they still do not change any verdict.

8 Objections

8.1 Land valuation and the transition shock

The Georgist programme has always been vulnerable at the point of assessment. Separating land value from improvement value requires either a functioning market in vacant land, which is thin in dense areas, or a residual method that subtracts depreciated replacement cost from total property value, which is sensitive to assumptions. Denmark’s experience is instructive and not encouraging: its land assessment system required a lengthy suspension and rebuild after systematic valuation errors, and Denmark is a small, well-administered country with unusually good property registers.

Machine valuation is a genuine answer to part of this. The information required to estimate land value is largely spatial, transactional and public, and automated valuation at parcel level is now considerably better than the periodic manual assessment it would replace. Paper IV’s identity and registry layer would make the assessments auditable and appealable in a way that current systems are not. This is one of the places where the agentic infrastructure genuinely improves an old proposal.

The transition shock is harder. A land value tax is capitalised into land prices: a permanent annual tax of τ_L on land value reduces that value by approximately $\tau_L/(i + \tau_L)$, where i is the discount rate. At a 3 per cent tax and a 4 per cent discount rate, land values fall by about 43 per cent on announcement. The loss falls on current owners, who in most countries are leveraged households, and it falls disproportionately on those who bought most recently at the highest prices. This is a one-time confiscation of the capitalised value of future rents from a group selected mainly by the timing of their purchase, and no amount of theoretical elegance about the neutrality of land taxes addresses it. Feldstein’s analysis of the incidence of a tax on pure rent showed that in general equilibrium the burden does not stay where the static analysis puts it, because the tax alters portfolio composition and hence the return on other assets [13].

The available mitigations are slow phase-in over decades, which sacrifices most of the revenue in the interim, or compensation, which sacrifices the rest. Oates and Schwab’s study of Pittsburgh’s graded tax found real effects on building activity but at rates far below those contemplated here [42]. We do not regard this objection as answered.

8.2 Consumption taxes and distribution

The idea of taxing what people take out of the economy rather than what they put in is old and respectable; Kaldor made the case for an expenditure base seventy years ago [27], and the subsequent global spread of the value added tax is one of the more consequential fiscal developments of the past half century [28]. The distributional objection is equally old.

A consumption tax is regressive against annual income and roughly proportional against lifetime income, since all income is eventually consumed or bequeathed. The standard response is that lifetime incidence is the relevant measure, and there is something to it, but it is weaker than usually claimed. Bequests break the lifetime equivalence at the top, where the largest fortunes are never consumed by their owners at all. Liquidity constraints break it at the bottom, where annual incidence is what determines whether a household can pay rent this month.

The architecture’s own answer is the basic consumption exemption or rebate: a universal transfer sized to the tax on a subsistence consumption bundle, which converts a proportional consumption tax into a progressive one at the bottom while leaving marginal rates untouched. This works arithmetically, and it costs revenue. Table 8 prices that cost rather than leaving it as a caveat. A rebate covering the first 20 per cent of average consumption raises the required American rate from 37.2 to 46.6 per cent tax-exclusive, and the Danish from 138.1 to 172.6. The distributive property is worth having and it is not free, and both halves of that belong in the same sentence as the rate.

There is a deeper point that the framework should concede. The income tax is not merely a revenue instrument; it is the principal mechanism through which advanced states adjust the distribution of market outcomes. Diamond and Saez’s case for progressive taxation rests on social welfare weights and empirically estimated elasticities, not on administrative convenience [10]. Abolishing progressive income taxation removes that instrument, and the proposed replacements are at best distributionally neutral. If the agentic economy concentrates returns on capital and autonomous systems, which Paper II suggests it may, then removing the redistributive instrument at the same moment that market inequality rises is a choice with consequences the framework has not yet priced. The endowment plus citizen dividend model, on the Alaska pattern, is the natural response. The Alaska Permanent Fund stood at roughly \$91 billion in mid-2026 and pays every resident an annual dividend [1], and Jones and Marinescu find no adverse aggregate employment effects from that programme [25]. A dividend funded from a fund that does not yet exist is not, however, a distributional policy.

There is a further point that Section 2 sharpened. The lifetime identity, not the Atkinson–Stiglitz theorem, is what makes a uniform consumption tax and a flat labour income tax close to equivalent. What the proposal actually gives up, then, is the *nonlinearity* of the schedule rather than the taxation of labour as such: the ability to set marginal rates that rise with ability, which is the entire content of the optimal income tax literature from Mirrlees onward

[37, 50]. A consumption tax with a basic rebate recovers some progressivity at the bottom and none at the top. If the framework wants a progressive outcome it will have to obtain it from the asset side, through the ownership of the endowment and the distribution of its returns, rather than from the tax schedule.

8.3 The political economy of a very large state endowment

Suppose the accumulation succeeded and a state came to own financial assets worth several times GDP. Three problems follow, and the framework's own constitutional commitments make them sharper rather than softer.

The first is governance. A fund of that size is a standing temptation, and the political economy literature on resource funds is largely a catalogue of raids. Norway's insulation rests on a fiscal rule, a professional manager at arm's length from the treasury, and a political consensus that has held for thirty years. Singapore's rests on constitutional protection through the elected presidency. Neither is a technical fix; both are contingent political settlements that a determined majority could undo. The proposal requires such a settlement to hold not for thirty years but for a century of accumulation followed by indefinite maintenance.

The second is market distortion. Norway's fund holds roughly 1.5 per cent of all listed equities globally and has become a significant governance actor by default. A fund of seven times American GDP would hold a controlling or near-controlling position in world capital markets. A state that owns the majority of the productive capital stock is exercising economic power whether or not it intends to, and Article VIII of the constitutional framework in Paper V, which forbids any single entity from obtaining unrestricted control of the economic protocol, applies to the state's own portfolio as much as to any corporation.

The third is the one the framework is least equipped for. If the state's revenue comes from a diversified claim on global production, the state's fiscal interest becomes aligned with asset prices rather than with domestic wages. That is a significant realignment of whose prosperity the government has a budgetary stake in, and it operates in exactly the opposite direction from the political economy that produced the modern welfare state.

8.4 If taxes shrink, what anchors the currency?

An objection from the other side of this series, and one that Paper I makes unavoidable. The chartalist tradition holds that state money is accepted because the state accepts it in payment of taxes [30]. Modern money theory sharpens this into a claim about sequencing: government spending creates the obligations that taxation subsequently extinguishes, so the tax liability is what generates the demand for the state's liability in the first place [63]. Minsky's version of the point is that anyone can create money and the difficulty is getting it accepted [34], with

the state enjoying a privileged position precisely because it can compel acceptance through the tax system.

If that is right, driving the taxes in (1) toward zero looks like it removes the foundation of the currency along with the tax. A state that taxes nothing has nothing to compel the use of its money, and Paper I's residual settlement asset loses its special status.

The objection has less force than it first appears, and the reason is specific. What the chartalist mechanism requires is a recurring, enforceable, universal liability denominated in the state's unit of account. It does not require that liability to fall on income. Every pillar in (2) generates exactly such a liability, and one of them generates a better one than the income tax does. A land value tax is levied annually on an asset that cannot be hidden, moved, or made to disappear by declining to work, and it falls on holders whether or not they transact in any given year. As a device for creating a standing, inescapable demand for the state's money it is strictly superior to a tax that can be avoided by earning nothing. Henry George and Knapp turn out to be complementary rather than opposed, which is not a connection either of them drew.

The residual worry is about scale rather than mechanism. A state collecting 8 per cent of GDP generates less demand for its unit of account than one collecting 45, and Paper I's analysis suggests the quantity of settlement assets required is set by the residual of the obligation graph rather than by the tax take. We regard this as unresolved. It is the principal question this paper leaves for Paper IV, which specifies the settlement layer.

8.5 Income survives because it is the broadest feasible base

This is the strongest objection and we take it last because the paper's own evidence partly supports it.

The case runs as follows. Income is taxed not because anyone thinks it the ideal base but because it is where the money is. Labour compensation is the majority of national income, it is observable through third-party reporting, and the enforcement technology is mature. Every alternative base proposed here is smaller, and the ones that are not smaller do not exist yet. On this reading the income tax is the equilibrium outcome of a constrained optimisation, and the constraint is real.

Our reply has two parts and the second is a concession.

The first part is that the constraint is technological, and the argument of Section 2 is that the technology is changing. Kleven, Kreiner and Saez's mechanism depends on production being organised inside firms with employees whose records are hard to falsify [29]. As production moves toward autonomous agents transacting continuously and netting obligations, the quantity being reported becomes ambiguous even when reporting itself becomes perfect. Perfect visibility

into an ill-defined quantity is not a tax base.

The second part is the concession. Nothing in this paper's arithmetic shows that the alternative bases can carry a modern state, and most of it shows the opposite. If the income tax base erodes for the reasons we describe, the honest conclusion is not that the alternatives will step in. It is that states will face a genuine financing crisis and will respond by taxing something, most likely by pushing consumption taxation to rates that this paper has just shown to be near the edge of feasibility, or by taxing the agentic production the framework wants to leave alone. Predicting the erosion of a tax base is not the same as having a replacement for it, and this paper should not be read as claiming otherwise.

9 A Transition Path

The programme is available in one step for the United States and Singapore, not at all for Denmark, and not on any near horizon for Norway, whose reclassified requirement is still 98.9 per cent tax-exclusive on an unverified consumption base. That leaves the question of what a government should actually do first. The sequence below is ordered by what the arithmetic supports rather than by what is easiest to legislate, and it holds whether or not a country ever reaches full abolition.

Build the consumption base before raising the rate. Everything in Section 6 turns on κ , and the difference between New Zealand's 0.96 and the OECD's 0.58 is worth more than any plausible land tax. A country that broadens its base to a single rate with no exemptions has done most of the work of abolition before touching the income tax at all; a country that raises the rate on a leaky base has taxed its remaining base twice and collected less than it expected. This ordering is the opposite of the political path of least resistance, which is to keep the exemptions and raise the rate.

Then take the taxes whose replacement is arithmetically available. In the United States and Singapore the corporate income tax is the smallest component of the productive-income aggregate, at 2.20 and 4.41 per cent of GDP. That is not general: in Denmark it exceeds social contributions, which are near zero, and in Norway it is the *largest* component at 12.00 per cent of GDP, three quarters of which is petroleum rent rather than a tax on domestic profit. The case for taking it first is therefore strongest precisely where it is smallest, and in Norway the question is not whether to abolish it but how much of it was ever a production tax. It is also the one where the Diamond–Mirrlees argument bites hardest, since it is a tax on intermediate production, and the one where the base is most mobile and therefore most expensively defended. There is a trap here that the efficiency argument itself sets, and stepping around it changes the recommendation. Production efficiency in Diamond–Mirrlees requires that pure profit be fully taxed; if it is not, the theorem's conditions fail and taxes on production can re-enter

the second-best optimum [11]. Simply abolishing the corporate income tax leaves pure profit untaxed and therefore cannot be justified by appeal to the very theorem invoked to justify it. The consistent move is to replace the corporate income tax with a tax on economic rent: a cash-flow tax, or an allowance for corporate equity that exempts the normal return while taxing the surplus above it. That instrument satisfies the theorem’s condition, taxes exactly what Axiom 1 does not protect, and leaves the marginal investment decision undistorted. It is also, in substance, what the Mirrlees Review recommended after the most thorough review of a national tax system yet undertaken [35, 36].

The distinction matters for the framework as a whole. Axiom 1 says production shall not be a taxable event. It does not say rent shall not be, and the entire revenue architecture in (2) is built on taxing rent. Corporate pure profit is rent. Exempting it would be a misreading of the axiom rather than an application of it. The intergenerational consequences of a shift of this kind are the subject of a large simulation literature that we do not attempt to replicate [5].

Next, price the externalities properly, and do it for allocative reasons rather than revenue ones, so that the shrinking base is a feature rather than a disappointment.

Third, build the land tax slowly. Announce a schedule that phases in over twenty to thirty years so that capitalisation is spread across successive owners rather than imposed on current ones, accept that the revenue arrives late, and use the agentic valuation and appeal infrastructure of Paper IV to make the assessments defensible.

Fourth, and this is the part that has to be argued honestly rather than sold, begin the endowment. Contributions have to come from somewhere and there is no painless source. Candidates that do not violate Axiom 1 include the resource rents themselves, the proceeds of spectrum and concession auctions, the land tax revenue in excess of current spending, and any public equity stake taken in exchange for public investment in compute, infrastructure or research. That last category is the most promising and the least developed. A state that funds foundational research or underwrites energy infrastructure for autonomous production has a legitimate claim on a share of the resulting returns, and taking that share as equity at the point of investment is consistent with the axiom in a way that taxing the resulting profits later would not be.

The distinction should not be oversold, and a referee pressed us on exactly how far it can be taken. In general equilibrium the two cash flows are close to equivalent: a dividend on state-held equity and a corporate income tax both transfer a share of the returns to production into public hands, either can distort investment depending on how it is structured, and if the state funds its equity purchases by taxing current consumption then the household lifetime budget constraint need not improve at all. What separates them is timing and contingency: an equity stake is priced when the investment is made and its return is contingent on success, whereas a profits tax is imposed after the fact on an investment made under different expectations. That

difference is real, and it is the same difference that makes a royalty and an auctioned concession different instruments, but it is a difference of degree.

The endowment will not be large enough to abolish personal income taxation within the lifetime of anyone now voting for it. Saying so plainly is a better foundation for the proposal than the alternative, because the alternative is a promise that the arithmetic in Section 5 does not support.

10 Limitations and Falsification

Several limitations bear on how much weight the results can carry.

The revenue-neutrality assumption holds expenditure fixed. If the agentic economy reduces the real cost of publicly provided services substantially, which is plausible for administration and possible for health and education delivery, the required G falls and every gap in Table 5 narrows proportionally. A 20 per cent reduction in the real cost of public provision would cut the American endowment requirement from 7.40 to 5.9 times GDP. We have not modelled this because we have no defensible basis for choosing the magnitude, but it is the single assumption most likely to be wrong in the favourable direction.

The composition figures are drawn from different vintages and, in the Norwegian case, do not reconcile. We have handled the Norwegian problem by reporting a range, and the Norwegian verdict does not depend on where in that range the truth lies, but the figure should be corrected against primary OECD tables before any of these numbers are used for anything other than the comparative purpose they serve here. Several other quantities are approximations and are marked in place: the Danish property tax composition, the UAE VAT share of GDP, and all sovereign fund estimates for the UAE, which are third-party figures for entities that do not report.

The return assumptions are conventional and could be wrong in either direction. Everything in Proposition 1 scales as $1/\rho$, so an endowment earning 6 per cent sustainably rather than 2.5 needs to be 58 per cent smaller. We regard a sustained real draw above 3 per cent net of growth as implausible for a fund large enough to be a significant share of world capital, since such a fund earns the market return by construction, but this is a judgement rather than a result.

The partial equilibrium assumption, and why it flatters the proposal. We treat r as a parameter, and this is the most consequential simplification in the paper. It is also not neutral: relaxing it makes the results worse rather than better, which is why we flag it here rather than leaving it to a reader to discover.

A public endowment of seven to twelve times GDP is not a price taker. In the United States and Denmark such a fund would exceed the entire private net worth of the country, and

accumulating it means adding that much capital to world markets. Standard capital theory says the marginal product of capital falls as the stock rises, so the act of accumulating the endowment depresses the return the endowment earns. Under the stewardship rule the damage is amplified, because the sustainable draw is $r - g$ rather than r , and a fall in r toward g shrinks the draw much faster than it shrinks the return. The requirement $w^* = \gamma/(r - g)$ then diverges hyperbolically:

$r - g$	3.5%	3.0%	2.5%	2.0%	1.5%	1.0%
w^* (US, $\gamma = 18.51$)	5.29	6.17	7.40	9.26	12.34	18.51

A one point fall in the sustainable draw, from 2.5 to 1.5 per cent, raises the American requirement from 7.40 to 12.34 times GDP. Since the accumulation itself is what pushes r down, there is a fixed point somewhere, and it is not obvious that a fixed point with a finite endowment exists for the larger targets. Establishing whether it does requires a general equilibrium model with an endogenous return, which we have not built. We note the direction of the bias and decline to claim more: every number in Table 5 should be read as a lower bound on the true requirement.

Finally, the whole exercise is conditional on a transformation of production that has not occurred. If autonomous agents do not come to occupy a large share of production, or if they do so inside conventional corporate structures that report conventional income, then the administrative argument of Section 2 does not apply and the income tax continues to work exactly as it does now.

The mobile-rent problem. One limitation deserves separating out because it attacks the architecture rather than the estimates. Section 2.3 argued that if agent capability is freely replicable, the surplus settles on inelastic inputs that the five pillars already reach. Frontier model access is not freely replicable. It is distributed under contract and protected by secrecy, and to the extent that holds, the rent accrues to the model provider rather than to energy, land or spectrum.

That rent has the opposite properties to the ones the architecture relies on. It is intangible, it has no physical situs, its transfer prices are set inside a single firm, and it can be booked in whichever jurisdiction offers the best treatment. A tax system built on immobile bases has no instrument that reaches it, and the international experience of trying to tax exactly this kind of return, through two decades of transfer pricing reform and most recently a global minimum rate, is not encouraging.

The uncomfortable implication is that the pillar architecture works best in the world where AI capability is commoditised and worst in the world where a few providers hold durable market power. That is the opposite of the ordering most people expect, since the concentrated world is

the one usually described as most in need of redistribution. We do not have a solution. The honest statement is that a residual tax on provider rent, which is a tax on profit and therefore on production, may turn out to be the one production tax the architecture cannot dispense with.

What would falsify this. Three observations would count against the paper's central claims. If a country succeeded in raising more than about 5 per cent of GDP from a pure land value tax on a sustained basis, the rent ceiling result would be wrong and the Georgist pillar would be far more capable than we have allowed. If a sovereign fund delivered sustained real returns net of growth above 4 per cent at a scale of several times GDP, the endowment requirements would fall by a third or more and the accumulation times with them. And if a large economy demonstrated that agentic production could be taxed on a throughput or settlement basis at low rates without measurable distortion, the entire premise that income becomes unadministrable would need revisiting, since a settlement tax would be a broader base than anything considered here.

11 Conclusion

This paper has argued for the abolition of taxes on labour and productive income, and has priced it. The two halves belong together, because the argument is common and the price is usually left out.

The case for abolition is that income taxation reaches wage earners with a completeness it never achieves against income that can be retained, deferred or converted, and that under agentic production the base itself loses its referent. Neither claim requires anyone to have intended the outcome. Both follow from the enforcement technology.

The price is a broad tax on final consumption, and for two of the five countries here it is payable. The United States can abolish every tax on earning at 37.2 per cent tax-exclusive, 27.1 per cent tax-inclusive, on a base with New Zealand's coverage; 46.6 per cent with a universal rebate that makes the effective schedule progressive in consumption. Singapore can do it at 35.3 per cent, or from investment returns alone with its consumption tax rising by a third of a point. These are unprecedented rates on an unprecedented base, and they are not absurd ones.

For Denmark it is not payable. The requirement is 138 per cent tax-exclusive on the household base and 78 per cent on the broadest base its own value added tax reaches, and no enforcement regime recovers a base at those rates on an open border. The cause is not Danish administration, which is excellent, but Danish household consumption at 43.6 per cent of GDP against America's 67.9. A state that routes lifetime consumption through its own budget has spent the base it would need to tax. Norway sits between, closer to Denmark than to the

United States once petroleum rent is set aside.

The scarcity pillars are supplements and this paper has been more definite about that than the tradition it draws on. No country has ever raised more than 1.09 per cent of GDP from a pure land tax or 0.44 per cent from carbon. Full confiscation of American land rent covers between a fifth and a third of the American gap. These are good taxes and small ones, and presenting them as the replacement for eighteen points of GDP is not a policy.

The public capital endowment lowers the rate and does not eliminate it. Abolition is available at the start rather than after a century of saving, which is a real correction to the framing this argument usually receives: each times-GDP of accumulated public capital drawn at 4 per cent removes four points of GDP from what consumption must raise, taking the American rate from 37.2 to 30.7 at one times GDP and to 24.2 at two. But a fund drawn on while it accumulates converges rather than compounding. At a 5 per cent contribution and a 4 per cent draw it converges to 3.33 times GDP, the consumption rate converges to 23.6 per cent tax-exclusive, and zero is not reached at any date. Lower draws raise the ceiling and slow the approach in the same proportion, pushing the zero out to several centuries. The destination is a lower consumption tax, not the absence of one, and we corrected this paper's earlier claim to the contrary rather than leaving it standing.

What remains unresolved is stated in Section 10 and we do not want it lost in a summary. If frontier model capability stays proprietary rather than becoming commoditised, the surplus of an automated economy accrues to providers as mobile, situs-free rent that no immobile-base pillar reaches, and a residual tax on that rent may be the one tax on production the architecture cannot dispense with. The programme works best in the world where artificial intelligence is a commodity and worst in the world where it is a franchise.

Article VII of Paper V schedules the abolition of taxes on labour and productive income, names final consumption as the primary replacement, and requires a state to publish each year how much of its revenue it still raises from production and where that share is heading. The arithmetic in this paper is why the article is a schedule with a disclosure duty rather than a declaration, and why the emergency departure clause has to exist. It is also why the schedule points the right way.

A Derivations

A.1 The endowment ratio

Let $W_G(t)$ evolve in continuous time as

$$\dot{W}_G = rW_G + I_G - G,$$

the continuous analogue of (5). Write $w = W_G/Y$ with $\dot{Y}/Y = g$. Then

$$\dot{w} = \frac{\dot{W}_G}{Y} - w \frac{\dot{Y}}{Y} = (r - g)w + \iota - \eta,$$

where $\iota = I_G/Y$ and $\eta = G/Y$ are the contribution and withdrawal shares. Setting $\iota = 0$ and $\eta = \phi w$ recovers the draw-rule dynamics of Section 3: $\dot{w} = (r - g - \phi)w$, stationary at $\phi = r - g$. Requiring the withdrawal to finance a revenue share γ gives $\phi w = \gamma$, hence $w^* = \gamma/\phi$, which is (8).

A.2 Accumulation time

With $\eta = 0$ and constant $\iota = s$, the linear equation $\dot{w} = (r - g)w + s$ has solution

$$w(t) = \left(w_0 + \frac{s}{r - g} \right) e^{(r - g)t} - \frac{s}{r - g}.$$

Setting $w_0 = 0$ and solving $w(t^*) = w^*$ gives (9). The expression is increasing in w^* and decreasing in s , and it is defined for $r > g$; for $r \leq g$ the fund cannot reach a stationary positive ratio to GDP from contributions alone at any finite time, since the growth of the denominator outpaces compounding.

A.3 Land tax capitalisation

A parcel yielding perpetual rent R with discount rate i is worth $V = R/i$ untaxed. Under an annual tax at rate τ_L on assessed value, the owner's net receipt is $R - \tau_L V'$, where V' is the post-tax value, so $V' = (R - \tau_L V')/i$ and

$$V' = \frac{R}{i + \tau_L} = V \cdot \frac{i}{i + \tau_L}, \quad \frac{V - V'}{V} = \frac{\tau_L}{i + \tau_L}.$$

At $\tau_L = 3\%$ and $i = 4\%$ the capitalised loss is $3/7 \approx 43$ per cent, as quoted in Section 8.

A.4 Consumption tax yield

Let $c = C/Y$ be household final consumption net of the consumption taxes being replaced, as a share of GDP, and let $\kappa \in (0, 1]$ be the share of that base the tax actually reaches. Revenue as a share of GDP at tax-exclusive statutory rate τ_e is then $\tau_e c \kappa$, so the rate required to raise a share γ is $\tau_e = \gamma/(c\kappa)$, and the tax-inclusive equivalent is $\tau_i = \tau_e/(1 + \tau_e)$.

The coverage parameter κ is *not* the OECD's VAT Revenue Ratio. The revenue ratio is defined as actual revenue divided by the product of the standard rate and total final consumption net of the tax, where total final consumption includes government [43]. Because its denominator

differs from ours, a revenue ratio cannot be substituted for κ ; where we use 0.96 we are stipulating a ceiling on achievable household coverage by analogy with the best-performing system, not transferring a measurement.

B Data Notes

The fiscal aggregates used in Section 5 are reproduced below with their vintages and confidence markings. Every cell states its denominator explicitly, because the compiled table this study began from mixed shares of GDP with shares of total tax revenue across rows and that is how the Norwegian error survived as long as it did.

Country	Total	Composition used (all % of GDP)	Status
United States	25.63% of GDP (2023); 25.62% (2024)	PIT 10.25, SSC 6.03, property 2.91, G&S 4.25, CIT 2.20, payroll 0.03	Verified, OECD RS2025
Denmark	44.04% of GDP (2023); 45.18% (2024)	PIT 25.17, G&S 12.41 (VAT 9.14), CIT 3.69, property 1.72, SSC 0.08	Verified. <i>Grundskyld</i> % of GDP computed from DKK amounts; 2024–25 bills provisional
Norway	41.63% of GDP (2023); 40.19% (2024)	PIT 9.60, CIT 12.00, SSC 9.23, G&S 9.58 (VAT 7.40), property 1.17, payroll 0.05	Verified. Main-land ratio 41.3% net of petroleum; the widely quoted 42.8% is not reproducible
Singapore	19.80% of GDP incl. NIRC (FY2025 revised)	CIT 4.41, NIRC 3.44, GST 2.66, PIT 2.58; expenditure 17.91	Verified from MOF; GDP implied from the published revenue ratio

Country	Total	Composition used (all % of GDP)	Status
UAE	28.5% of GDP general govern- ment revenue (2023); 27.8% (2024)	No PIT; hydrocarbon (2023), 14.8 (2024); VAT 5% and excise combined only	No primary CIT yield exists. Sovereign fund sizes are external estimates except Mubadala

Bibliography. Twenty-five of the sixty-three references here were checked against full Crossref records rather than against memory or secondary listings. The check found two errors in our own draft. One reference did not exist: we had cited a “Global Carbon Budget 2025”, but the edition published in 2025 is the Global Carbon Budget 2024. The other was a wrong title on a real paper. We had recorded Sandmo (1975) as “Optimality Rules for the Taxation of Externalities” against a published title of “Optimal Taxation in the Presence of Externalities”, with volume, issue, page and year all correct. The second is the more dangerous of the two. An incorrect page range still lands a reader in the right article, whereas an incorrect title does not resolve at all and is invisible to every check short of attempting retrieval.

Two artefacts of the method should be recorded for anyone repeating it. A query that returns no journal record is not evidence that none exists: for seven of our references the working paper or preprint outranked the version of record, which appeared only on a direct lookup by identifier. And many older records store a start page alone, which should not be mistaken for a corrected page range. Where a closing page could not be confirmed we have followed the published convention and said so in the entry. Diamond and Mirrlees (1971) predates the identifier system and is unverifiable by this route; we cite Part II as 61(3) and note that some sources give 61(2).

Figures deliberately not used. A Norwegian mainland tax ratio of 42.8 per cent; a Norwegian personal income tax of 19.5 per cent of GDP (this is category 1000, personal and corporate combined); a United States total tax ratio of 25.2 per cent; a Danish 2023 total of 43.4 per cent (superseded by the 2025-edition revision to 44.04); a UAE hydrocarbon share of 40 per cent of revenue; and any yield figure for the UAE corporate income tax. Each appeared in the compiled source this study started from and each failed verification against primary tables [45, 53, 54].

References

- [1] Alaska Permanent Fund Corporation. Mid-fiscal-year review, fiscal year 2026. Technical report, Alaska Permanent Fund Corporation, Juneau, AK, 2026.
- [2] Iñaki Aldasoro and Ajit Desai. AI agents for cash management in payment systems. BIS Working Paper 1310, Bank for International Settlements, Basel, 2025.
- [3] Richard J. Arnott and Joseph E. Stiglitz. Aggregate land rents, expenditure on public goods, and optimal city size. *The Quarterly Journal of Economics*, 93(4):471–500, 1979. doi: 10.2307/1884466. Crossref reports start page 471 only; the closing page follows the published convention.
- [4] Anthony B. Atkinson and Joseph E. Stiglitz. The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1–2):55–75, 1976. doi: 10.1016/0047-2727(76)90041-4.
- [5] Alan J. Auerbach and Laurence J. Kotlikoff. *Dynamic Fiscal Policy*. Cambridge University Press, Cambridge, 1987.
- [6] Board of Governors of the Federal Reserve System. Financial accounts of the united states (z.1), first quarter 2026. Federal Reserve Statistical Release Z.1, Washington, DC, 2026. Released 11 June 2026.
- [7] Christophe Chamley. Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica*, 54(3):607–622, 1986. doi: 10.2307/1911310. Crossref reports the start page only; the closing page follows the published convention.
- [8] R. H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937. doi: 10.1111/j.1468-0335.1937.tb00002.x.
- [9] Morris A. Davis and Jonathan Heathcote. The price and quantity of residential land in the united states. *Journal of Monetary Economics*, 54(8):2595–2620, 2007. doi: 10.1016/j.jmoneco.2007.06.023.
- [10] Peter Diamond and Emmanuel Saez. The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, 25(4):165–190, 2011. doi: 10.1257/jep.25.4.165.
- [11] Peter A. Diamond and James A. Mirrlees. Optimal taxation and public production i: Production efficiency. *The American Economic Review*, 61(1):8–27, 1971. Pre-DOI American Economic Review issue; not indexed in Crossref. Volume, issue and pages follow the published convention.

- [12] Peter A. Diamond and James A. Mirrlees. Optimal taxation and public production ii: Tax rules. *The American Economic Review*, 61(3):261–278, 1971. Pre-DOI American Economic Review issue; not indexed in Crossref. Cited as 61(3); some secondary sources give 61(2).
- [13] Martin Feldstein. The surprising incidence of a tax on pure rent: A new answer to an old question. *Journal of Political Economy*, 85(2):349–360, 1977. doi: 10.1086/260566.
- [14] William G. Gale. The required tax rate in a national retail sales tax. *National Tax Journal*, 52(3):443–457, 1999. doi: 10.1086/ntj41789734.
- [15] Henry George. *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions and of Increase of Want with Increase of Wealth*. D. Appleton and Company, New York, 1879.
- [16] Roger Gordon and Wei Li. Tax structures in developing countries: Many puzzles and a possible explanation. *Journal of Public Economics*, 93(7–8):855–866, 2009. doi: 10.1016/j.jpubeco.2009.04.001.
- [17] David Graeber. *Debt: The First 5,000 Years*. Melville House, Brooklyn, NY, 2011.
- [18] Gillian K. Hadfield and Andrew Koh. An economy of AI agents. In *The Economics of Transformative AI*. University of Chicago Press for the National Bureau of Economic Research, Chicago, 2025. Also available as arXiv:2509.01063.
- [19] Inland Revenue Authority of Singapore. Annual report fy2024/25. Technical report, Inland Revenue Authority of Singapore, Singapore, 2025.
- [20] A. Mitchell Innes. What is money? *The Banking Law Journal*, 30:377–408, 1913.
- [21] International Energy Agency. *CO2 Emissions in 2024*. International Energy Agency, Paris, 2025. Institutional statistical source; not Crossref-indexed, cited by convention.
- [22] International Monetary Fund. Norway: 2025 article iv consultation—press release; staff report; and statement by the executive director for norway. IMF Country Report 25/249, International Monetary Fund, Washington, DC, 2025.
- [23] International Monetary Fund. United arab emirates: 2025 article iv consultation—press release; staff report; and statement by the executive director for the united arab emirates. IMF Country Report 25/327, International Monetary Fund, Washington, DC, 2025.
- [24] Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976. doi: 10.1016/0304-405X(76)90026-X.

- [25] Damon Jones and Ioana Marinescu. The labor market impacts of universal and permanent cash transfers: Evidence from the alaska permanent fund. *American Economic Journal: Economic Policy*, 14(2):315–340, 2022. doi: 10.1257/pol.20190299.
- [26] Kenneth L. Judd. Redistributive taxation in a simple perfect foresight model. *Journal of Public Economics*, 28(1):59–83, 1985. doi: 10.1016/0047-2727(85)90020-9.
- [27] Nicholas Kaldor. *An Expenditure Tax*. George Allen and Unwin, London, 1955.
- [28] Michael Keen and Ben Lockwood. The value added tax: Its causes and consequences. *Journal of Development Economics*, 92(2):138–151, 2010. doi: 10.1016/j.jdeveco.2009.01.012.
- [29] Henrik Jacobsen Kleven, Claus Thustrup Kreiner, and Emmanuel Saez. Why can modern governments tax so much? an agency model of firms as fiscal intermediaries. *Economica*, 83(330):219–246, 2016. doi: 10.1111/ecca.12182.
- [30] Georg Friedrich Knapp. *The State Theory of Money*. Macmillan, London, 1924. Abridged English translation of *Staatliche Theorie des Geldes*, Leipzig, 1905.
- [31] William D. Larson. New estimates of value of land of the united states. Bea working paper, Bureau of Economic Analysis, U.S. Department of Commerce, Washington, DC, 2015.
- [32] Perry Mehrling. *The New Lombard Street: How the Fed Became the Dealer of Last Resort*. Princeton University Press, Princeton, NJ, 2011.
- [33] Ministry of Finance, Singapore. Analysis of revenue and expenditure, financial year 2025. Technical report, Ministry of Finance, Republic of Singapore, Singapore, 2025.
- [34] Hyman P. Minsky. *Stabilizing an Unstable Economy*. Yale University Press, New Haven, CT, 1986.
- [35] James Mirrlees, Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba, editors. *Dimensions of Tax Design: The Mirrlees Review*. Oxford University Press for the Institute for Fiscal Studies, Oxford, 2010.
- [36] James Mirrlees, Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba. *Tax by Design: The Mirrlees Review*. Oxford University Press for the Institute for Fiscal Studies, Oxford, 2011.
- [37] James A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971. doi: 10.2307/2296779. Crossref reports the start page only; the closing page follows the published convention.

- [38] Matthew N. Murray. Would tax evasion and tax avoidance undermine a national retail sales tax? *National Tax Journal*, 50(1):167–182, 1997. doi: 10.1086/ntj41789248.
- [39] William D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, 2017. doi: 10.1073/pnas.1609244114.
- [40] Norges Bank Investment Management. Government pension fund global: Annual report. Technical report, Norges Bank Investment Management, Oslo, 2025.
- [41] Norwegian Ministry of Finance. Meld. st. 22 (2024–2025): Long-term perspectives on the norwegian economy and the management of the government pension fund. Technical report, Norwegian Ministry of Finance, Oslo, 2025.
- [42] Wallace E. Oates and Robert M. Schwab. The impact of urban land taxation: The pittsburgh experience. *National Tax Journal*, 50(1):1–21, 1997. doi: 10.1086/ntj41789240.
- [43] Organisation for Economic Co-operation and Development. *Consumption Tax Trends 2024: VAT/GST and Excise, Core Design Features and Trends*. OECD Publishing, Paris, 2024.
- [44] Organisation for Economic Co-operation and Development. *National Accounts of OECD Countries: Financial Balance Sheets and Non-Financial Assets*. OECD Publishing, Paris, 2025. Institutional statistical source; not Crossref-indexed, cited by convention.
- [45] Organisation for Economic Co-operation and Development. *Revenue Statistics 2025: Tax Revenue Trends in the OECD*. OECD Publishing, Paris, 2025.
- [46] Organisation for Economic Co-operation and Development. *Revenue Statistics in Asia and the Pacific 2025*. OECD Publishing, Paris, 2025.
- [47] Thomas Piketty and Gabriel Zucman. Capital is back: Wealth-income ratios in rich countries 1700–2010. *The Quarterly Journal of Economics*, 129(3):1255–1310, 2014. doi: 10.1093/qje/qju018.
- [48] Frank P. Ramsey. A contribution to the theory of taxation. *The Economic Journal*, 37 (145):47–61, 1927. doi: 10.2307/2222721. Crossref reports the start page only; the closing page follows the published convention.
- [49] Republic of Singapore. Constitution of the republic of singapore, article 142. Reprint, Singapore Statutes Online, 2016. Net Investment Returns framework; Art. 142(1A) and 142(4).
- [50] Emmanuel Saez. Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, 68(1):205–229, 2001. doi: 10.1111/1467-937X.00166.

- [51] Agnar Sandmo. Optimal taxation in the presence of externalities. *The Swedish Journal of Economics*, 77(1):86–98, 1975. doi: 10.2307/3439329.
- [52] Joel Slemrod. Optimal taxation and optimal tax systems. *Journal of Economic Perspectives*, 4(1):157–178, 1990. doi: 10.1257/jep.4.1.157.
- [53] Statistics Denmark. Statbank SKAT: Taxes and duties by type. Statistics Denmark Statbank, Copenhagen, 2026. Updated 3 June 2026.
- [54] Statistics Norway. National accounts and government revenue statistics. Statistisk sentralbyrå, Oslo, 2026. Institutional statistical source; not Crossref-indexed, cited by convention.
- [55] Ludwig Straub and Iván Werning. Positive long-run capital taxation: Chamley–judd revisited. *The American Economic Review*, 110(1):86–119, 2020. doi: 10.1257/aer.20150210.
- [56] T. Nicolaus Tideman. A tax on land value is neutral. *National Tax Journal*, 35(1):109–111, 1982. doi: 10.1086/ntj41862425.
- [57] Nenad Tomasev, Matija Franklin, Joel Z. Leibo, Julian Jacobs, William A. Cunningham, Iason Gabriel, and Simon Osindero. Virtual agent economies. Technical Report arXiv:2509.10147, arXiv, 2025.
- [58] United States Congress. H.r. 25, the FairTax act of 2023, 118th congress, section 101(b)(1). U.S. Government Publishing Office, Washington, DC, 2023. Statutory text verified at govinfo.gov; the rate is stated as a percentage of “gross payments”, which the Act defines to include the tax itself, and is therefore tax-inclusive.
- [59] U.S. Environmental Protection Agency. Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. Technical report, U.S. Environmental Protection Agency, Washington, DC, 2023. Institutional statistical source; not Crossref-indexed, cited by convention.
- [60] Oliver E. Williamson. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press, New York, 1985.
- [61] World Bank. State and trends of carbon pricing 2025. World Bank, Washington, DC, 2025. Institutional statistical source; not Crossref-indexed, cited by convention.
- [62] World Bank. World development indicators: Households and NPISHs final consumption expenditure (% of GDP). World Bank, Washington, DC, 2026. Institutional statistical source; not Crossref-indexed, cited by convention.
- [63] L. Randall Wray. *Modern Money Theory: A Primer on Macroeconomics for Sovereign Monetary Systems*. Palgrave Macmillan, Basingstoke, 2 edition, 2015.