CIGS The Canon Institute for Global Studies

CIGS Working Paper Series No. 26-022E

How Large Are the Effects of Inheritance Taxation?

Lessons from Japan

Selahattin İmrohoroğlu (Marshall School of Business, University of Southern California, and CIGS,)
Fang Yang (Federal Reserve Bank of Dallas)

October 26, 2026

※Opinions expressed or implied in the CIGS Working Paper Series are solely those of the author, and do not necessarily represent the views of the CIGS or its sponsor.
※CIGS Working Paper Series is circulated in order to stimulate lively discussion and comments.
※Copyright belongs to the author(s) of each paper unless stated otherwise.

General Incorporated Foundation
The Canon Institute for Global Studies
一般財団法人 キヤノングローバル戦略研究所
Phone: +81-3-6213-0550 https://cigs.canon/

How Large Are the Effects of Inheritance Taxation? Lessons from Japan*

Selahattin İmrohoroğlu† and Fang Yang‡

October 26, 2026

Abstract

Japan taxes inheritances more heavily than almost any other advanced economy: revenue is about 0.5 percent of GDP and nearly 10 percent of decedents leave taxable estates. We study reform in an overlapping generations model with idiosyncratic earnings risk, voluntary and accidental bequests, and intergenerational transmission of earnings ability, calibrated to Japanese aggregates, earnings inequality, and the distribution of taxable estates. Because Japan is the world's largest net foreign creditor, we consider both a small open economy and a closed economy. Eliminating the tax raises household net worth by 15.3 percent in the small open economy, where the additional saving is held abroad, and by 1.8 percent in the closed economy, where domestic capital rises by 2.2 percent. An economy recalibrated without a voluntary bequest motive implies that about two-thirds of the benchmark bequest flow is accidental. Raising the exemption to the current U.S. level lowers the top one percent's wealth share in both economies but lowers ex-ante welfare in the closed economy, where fewer than one newborn in five gains. Under the 2070 mortality projections the responses are smaller throughout.

Keywords: Inheritance taxation, Bequests, Reform, Wealth Inequality, Japan
JEL Classification: D31; E21; H24; J14

*We are grateful to Mariacristina De Nardi for useful comments. We thank Benjamin Hoham for outstanding research assistance. This research was supported through computational resources provided by the Big-Tex High Performance Computing Group at the Federal Reserve Bank of Dallas. The opinions and conclusions are solely those of the authors and do not reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Marshall School of Business, University of Southern California and CIGS, selo@marshall.usc.edu.

‡Research Department, Federal Reserve Bank of Dallas, fang.yang@dal.frb.org.

1 Introduction

Inheritance taxation is controversial. Taxing inheritances heavily can reduce income and wealth inequality and weaken the transmission of economic advantage across generations, but it may also discourage the accumulation of capital. The quantitative literature that evaluates this trade-off has focused almost exclusively on the United States, where the estate tax is small: it raises little revenue, and only a tiny fraction of estates ever pay it. In such an environment, perhaps unsurprisingly, quantitative general equilibrium models find that even large changes in estate taxation have small aggregate effects; see, for example, Cagetti and De Nardi (2009) and De Nardi and Yang (2016).

Japan is different. Inheritance tax revenue amounts to 0.50% of GDP in Japan versus 0.088% in the U.S. today.[1] The tax base is also far broader: after the 2015 reduction of the basic exemption, roughly 10% of Japanese decedents leave taxable estates, compared with about 0.1% in the U.S. Finally, the Japanese statutory schedule is considerably more progressive, with marginal rates rising from 10% to 55%. Japan is thus an important laboratory for studying inheritance taxation, and reforms to it can produce significant macroeconomic, distributional, and welfare effects.

In this paper we quantify these effects. We build an incomplete-markets, overlapping generations model in the tradition of De Nardi (2004) and De Nardi and Yang (2016), in which households face idiosyncratic earnings risk and borrowing constraints, derive utility from leaving bequests, and transmit earnings ability to their children. Two intergenerational channels—voluntary and accidental bequests, and the inheritance of human capital—jointly generate a realistic concentration of wealth and of taxable estates. We adapt the framework to Japan by incorporating government debt held in household portfolios, the Japanese public pension system with its lump-sum basic pension and earnings-related component, and fiscal aggregates following Hansen and İmrohoroğlu (2023). We calibrate the model to Japanese data on earnings inequality (Kitao and Yamada 2025), the intergenerational persistence of earnings, and—using National Tax Agency statistics—the fraction of decedents paying inheritance taxes (9.9%), the ratio of inheritance tax revenue to output (0.50%), and the 90th percentile of the distribution of bequests. The implied effective exemption of ¥34.8 million is remarkably close to the statutory basic exemption, and the implied effective flat rate of 39.5% lies within the statutory progressive schedule.

[1] The U.S. figure refers to the current tax, whose exemption is $15 million per person (permanent from 2026), with a top statutory rate of 40% and roughly one decedent in a thousand leaving a taxable estate. The quantitative studies cited above model the tax as it stood around 2000: De Nardi and Yang (2016) calibrate an effective exemption of $756,000 and an effective marginal rate of 21.4%, reaching 1.9% of decedents and raising 0.33% of GDP. Against that tax, Japan's tax raises about one and a half times as much, rather than six times. Section 2 discusses the implications for our comparisons. On a basis fully harmonized across countries by the OECD—revenues from estate, inheritance, and gift taxes as a share of total tax revenue, inclusive of social security contributions and subnational taxes—Japan collects about 1.3 percent, one of only four OECD countries above one percent, compared with roughly 0.4 percent in the United States (OECD 2021). South Korea, Belgium, and France are the other three; measured relative to GDP rather than total tax revenue, they collect about 0.7% of their GDP from these taxes.

Using a similar model calibrated to the United States, De Nardi and Yang (2016) find that reducing the threshold for estate taxation and raising the effective tax rate to its then-statutory value of 55% reduces capital and output but also reduces wealth inequality, especially at the top end, with significant welfare gains for a newborn at the cost of very large losses for the very wealthiest. All of these effects, however, are quite small, reflecting the narrow base and low revenue of the U.S. estate tax even at that time. We keep the household side of our model deliberately close to theirs, so that differences in preferences, demographics and the tax are not confounded with differences in structure. Our model differs from theirs in two respects: government debt held in household portfolios, and the small open economy. Both are motivated by Japanese institutions, and both turn out to matter for the size of the aggregate response, which is itself one of our findings.

Because Japan has been the world's largest net foreign creditor for more than three decades, with a net international investment position on the order of 80% of GDP, we report two polar cases rather than privileging one. In the *small open economy*, which we take as the baseline, the return to capital r_k is pinned at the world rate, so the domestic capital stock, output and the wage are fixed and the difference between household net worth and the sum of domestic capital and government debt accumulates as net foreign assets. In the *closed economy*, which is the standard case in the literature and the one De Nardi and Yang (2016) use, that saving must be absorbed domestically and the interest rate falls. That case suits the literature's focus on the United States, which produces about a quarter of world output and is large enough to move the world interest rate; Japan, at about 4 percent, is not. Japan's small share of world output makes the small open economy the closer approximation; the closed economy shows how much changes if additional saving must be absorbed at home.

The difference between them is large enough to state first. Eliminating the inheritance tax raises household net worth by 15.3% at fixed prices. Since net worth is 3.94 times annual output, that is an increase equal to 60% of one year's output, and because the domestic capital stock is fixed by construction, all of it accrues as net foreign assets. In the closed economy the same reform raises the capital stock by 2.2% and output by 0.8%, lowering the return on household portfolios of capital and government debt from 3.76% to 3.64%. On a common household-wealth basis, the closed-economy increase is approximately 1.8%, making the small-open-economy response roughly nine times larger. The distributional effects are far more robust to the degree of capital mobility: the share of wealth held by the richest 1% rises from 17.4% to 18.7% at fixed prices and to 19.7% in the closed economy.

How do these responses compare with those for the United States? The closest benchmark is De Nardi and Yang (2016), whose closed-economy model shares our household structure. Their estate tax raises 0.33% of GDP from 1.9% of decedents; Japan's inheritance tax raises about one and a half times as much revenue from a share of decedents five times as large. In the comparable closed-economy experiment, abolishing the tax raises the capital stock about 1.4 times as much in Japan as in their model, close to the ratio of revenues. Because the two models are calibrated to different economies, this ratio

describes differences between them rather than isolating the role of any single feature.

Four further results are worth flagging. First, the choice of the instrument that absorbs the change in inheritance tax revenue matters far more than it does in the U.S. literature, where the tax is too small for the financing margin to be quantitatively relevant, and it reverses the ranking of the reforms. When the capital income tax adjusts instead of the labor income tax, eliminating the tax raises net worth by only 1.2% rather than 15.3%, because replacing a levy paid once per lifetime with an annual tax on the return to saving requires raising the capital income tax from 22.0% to 24.9%. The exemption reform is the exception: it induces enough additional accumulation that the capital income tax base expands by more than the revenue forgone, so that tax can *fall*, and the reform becomes the largest of the three. Second, the welfare verdict depends on whether capital markets are open. At fixed prices every newborn gains from the exemption reform, with an ex-ante gain of 0.061 times average income when the labor income tax adjusts; in the closed economy the same reform lowers ex-ante welfare by 0.022 times average income and only 18.8% of newborns gain, because the capital deepening it induces lowers the interest rate and that loss falls on precisely the households whose saving the reform was meant to encourage. Third, the size of the response depends on the bequest motive. An economy without a voluntary motive, recalibrated to the same wealth and inheritance-tax targets, generates about two-thirds of the benchmark bequest flow relative to wealth, and eliminating the tax raises its net worth by 6.5% rather than 15.3%. Life-cycle and precautionary saving alone thus produce substantial estates, but these are not unresponsive to the tax, which still changes heirs' receipts, fiscal financing and equilibrium returns. Fourth, replacing the survival probabilities with those projected for the 2070 birth cohort by the Institute of Population and Social Security Research shrinks almost every response, roughly by half in the baseline. Longer lives let households spend down their wealth, so aging lowers the accidental share of bequests from 67% to 58% under our model-based decomposition; on its own, this would make bequests more responsive to the tax. Responses shrink nonetheless because aging reduces the flow of bequests relative to wealth and delays inheritances.

The paper is organized as follows. Section 2 frames our contribution in the context of the previous literature. Section 3 presents the model and the two polar cases. Section 4 discusses the calibration. Section 5 highlights the implications of the calibrated model for the wealth distribution and for the importance of parental background. Section 6 studies the effects of reforming inheritance taxation in the small open economy, examines the role of the financing instrument, and reports the distributional effects. Section 7 reports the closed-economy results and compares them with De Nardi and Yang (2016). Section 8 evaluates welfare in both cases. Section 9 reports the alternative bequest specifications, including a model-based decomposition of the bequest flow into accidental and intentional components, and Section 10 studies population aging. Section 11 concludes.

2 Related Literature

Our paper builds on the quantitative literature on wealth inequality, intergenerational links, and estate taxation, and connects it to a growing literature on bequests and inheritance taxation in Japan.

Estate taxation, entrepreneurship, and wealth. Cagetti and De Nardi (2009) study estate taxation in an overlapping generations model with entrepreneurial ability, borrowing constraints, and intergenerational transfers, calibrated to the concentration of U.S. wealth. The tax matters little for small businesses but can distort investment and succession in large, illiquid firms, and eliminating it is not necessarily welfare-improving once the lost revenue is replaced by other taxes. We do not model entrepreneurship. Our earnings process includes a rare, extremely high state in the spirit of Castañeda, Díaz-Giménez, and Ríos-Rull (2003), which can be read as a reduced-form stand-in for successful entrepreneurs and other top earners, but it lasts a single year (Section 4), so large estates in our model come from windfalls that are saved rather than from business capital. Whether the tax falls on passive dynastic wealth or on constrained productive businesses is a margin our framework cannot evaluate.

Guo (2026) revisits inheritance and entrepreneurship in U.S. wealth inequality with occupational choice and a non-homothetic bequest motive. Disciplined by Survey of Consumer Finances evidence that more than half of the richest 1 percent report no inheritance at all, the model implies that inheritances play a limited role at the top of the U.S. distribution and that estate-tax reforms move wealth concentration less than commonly believed. In our Japanese calibration, by contrast, bequests account for the bulk of the advantage enjoyed by children of the richest households (Section 5).

Family background, bequests, and estate taxation. De Nardi and Yang (2016) combine two intergenerational channels—the transmission of earnings ability and the transmission of wealth through accidental and voluntary bequests—and show that voluntary bequests are needed to match the upper tail of the U.S. wealth distribution. In their model, higher estate taxation reduces concentration at the top with small aggregate effects, and raises ex-ante welfare by reducing the importance of being born into a wealthy family, at a large cost to children of rich families. Their view of the tax as a policy toward dynastic advantage, rather than mainly toward growth, is the lens of our results on parental background in Sections 5 and 6.

Because we compare our results with theirs throughout, it matters which version of the U.S. estate tax their model contains. Their benchmark is calibrated to an effective exemption of $756,000 and an effective marginal rate of 21.4%; it reaches 1.9% of decedents and raises 0.33% of GDP. Their main policy experiment raises the tax to the statutory levels then in force—a 55% rate above a $675,000 exemption—a reform toward a *broader* base and a *higher* rate.

The U.S. tax has since moved the other way. The exemption rose to $5 million under the 2010 and 2012 legislation, roughly doubled under the 2017 Tax Cuts and Jobs Act, and was set at $15 million per person, indexed and made permanent, by the 2025 reconciliation

act, effective from 2026. The top statutory rate is now 40%, roughly one decedent in a thousand leaves a taxable estate, and revenue is about 0.088% of GDP.

Two consequences follow. When we compare aggregate responses with theirs, we compare their tax as modeled with ours as modeled, not with current U.S. policy. And the exemption reform we study—raising x_b to ¥575 million, roughly 115 times average income—is calibrated to the *current* U.S. exemption, not to the one in their model. It is therefore a far more radical narrowing of the base than anything they consider, and asks what would happen if Japan adopted the U.S. tax *as it exists today*.

Our treatment of intergenerational links lies between two polar assumptions: pure life-cycle models in which bequests are only accidental, as in İmrohoroğlu, İmrohoroğlu, and Joines (1995), and models of two-sided altruism in which a family behaves as an infinitely-lived dynasty, as in Fuster, İmrohoroğlu, and İmrohoroğlu (2003), Fuster, İmrohoroğlu, and İmrohoroğlu (2007), and Fuster, İmrohoroğlu, and İmrohoroğlu (2008). As in De Nardi, İmrohoroğlu, and Sargent (1999) and De Nardi and Yang (2016), parents in our model value the bequests they leave and transmit earnings ability to their children, but each generation solves its own life-cycle problem. Where an economy sits on this spectrum matters for policy: Yang (2013) shows that the welfare gain from eliminating Social Security under impure altruism lies between those in the pure life-cycle and two-sided altruism models.

Bequests and inheritance taxation in Japan. Ono (2024) builds a life-cycle model for Japan in which bequests transmit physical capital and correlated productivity transmits human capital, following De Nardi (2004) and De Nardi and Yang (2016). Bequests matter in Japan because they transmit both resources and family advantage, and because Japanese inheritances consist largely of deposits and real estate rather than closely held businesses, the relevant policy margin may be less about entrepreneurship and more about whether inherited assets are held passively. Two pieces of Japanese evidence inform our bequest motive. Horioka (2014) finds that bequest motives in Japan are weaker and more selfish than in the United States, which argues against an altruistic, dynastic formulation in which parents value their children's utility and bequests respond to the children's own resources. Niimi and Horioka (2019) show that the Japanese elderly decumulate wealth very slowly, which a pure life-cycle model without a bequest motive cannot reproduce. Together they point to a bequest motive that is not altruistic, and we model it as a warm glow over the estate, a luxury good that keeps wealthy households holding assets late in life. The evidence supports this formulation rather than singling it out: slow decumulation is also consistent with precautionary saving for medical and long-term care. We rely on Kitao and Yamada (2025) for Japanese earnings and wealth inequality and on Hansen and İmrohoroğlu (2023) for the fiscal and national-accounts side of the economy.

Quantitative work on the Japanese tax is scarce despite its size, for reasons of data and institutions. Inheritance-tax returns are not available for research; household surveys underrepresent the top of the wealth distribution, which is the entire base of the tax; and liability depends on statutory heir shares and on preferential valuation of land and

unlisted business assets. Our approach, which summarizes the code by an effective rate and an effective exemption disciplined by National Tax Agency tabulations, is designed to work within these constraints.

Our contribution. Relative to these literatures, our contribution is to take the framework of De Nardi and Yang (2016) to the one advanced economy in which inheritance taxation is quantitatively important today, and to use it to isolate what governs the size of the aggregate response. Three features of the exercise are worth flagging here.

First, the Japanese data discipline the model tightly. The second-stage estimation recovers an effective exemption of ¥34.8 million—close to the statutory basic deduction of ¥30 million plus ¥6 million per heir—and an effective marginal rate of 39.5%, within the statutory range of 10 to 55 percent, even though no feature of the tax code is targeted. The model also matches the untargeted shares of wealth held by the top 5 to 40 percent of households and the life-cycle profiles of average net wealth and within-age-group inequality, including the slow decumulation of wealth in old age that motivates the bequest channel.

Second, we compare the benchmark with a separately recalibrated economy without a voluntary bequest motive, following Hendricks (2002). The no-motive economy generates roughly two-thirds of the benchmark bequest flow relative to wealth. This is a useful model-based comparison, but it does not separately identify the voluntary motive's contribution to reform responses or establish that the remaining flow is tax-unresponsive.

Third, we take seriously the fact that Japan is a large net foreign creditor. The standard closed-economy assumption forces additional household saving into the domestic capital stock and so damps the response through the interest rate. For an economy with a net international investment position near 80% of GDP, the opposite polar case—where the marginal unit of saving is placed abroad at an unchanged world return—is at least as defensible. We report both, and the gap between them turns out to be, together with the choice of financing instrument, one of the two largest sources of quantitative uncertainty in the paper.

3 The Model

The model is a discrete-time, incomplete-markets, overlapping-generations economy with an infinitely lived government.

3.1 Firm and Technology

There is a representative firm that rents capital and hires labor to produce goods using a Cobb-Douglas production function $K^{\alpha}L^{1-\alpha}$, where K is the aggregate capital stock, L is the aggregate labor input, and α is capital's share of output. The final goods can either be consumed or invested in physical capital, which depreciates at the rate δ.

3.2 The Government

We assume that the government makes exogenous purchases G and pays an exogenous interest rate r_d on outstanding debt D. We hold the government debt-output ratio fixed across stationary equilibria:

$$D = \bar{d}Y, \qquad Y = F(K, L),$$

where $\bar{d}$ is calibrated to the benchmark debt-output ratio. Thus, debt is unchanged in the small-open-economy inheritance-tax experiments, where production is fixed, but changes with output in the closed economy.

To finance these expenditures, the government taxes capital income at rate τ_k, labor income and public pension benefits at the rate τ_l, taxes interest earned on government debt at τ_d, and inheritances above a threshold x_b at the rate τ_b.

In addition, the government maintains a pay-as-you-go public pension system that pays benefits linked to an individual's realized average annual earnings $\widetilde{y}$, up to a cap $\widetilde{y}_c$, plus a lump-sum basic pension. These are financed through a payroll tax τ_s.

3.3 Demographics and Labor Earnings

Each model period is one year. Agents start their economic life at the age of 25 ($t = 1$). By age 35 ($t = 11$), the agents' children are born. The agents retire at age 67 ($t = 43$). From that period on, each household faces a positive probability of dying, given by $(1-p_t)$, which only depends on age.[2]

The maximum life span is age 94 ($T = 70$), and the population grows at a constant rate n. Figure A1 in Appendix A.1 displays the demographic structure of our overlapping generations model.

Total labor productivity of worker i at age t is given by $y_t^i = e^{z_t^i + \epsilon_t}$, in which ϵ_t is the deterministic age-efficiency profile. The process for the stochastic earnings shock z_t^i is: $z_t^i = \rho_z z_{t-1}^i + \mu_t^i,\ \mu_t^i \sim N(0, \sigma_\mu^2)$.

To capture the intergenerational correlation of earnings, we assume that the productivity of worker i at age 60 is transmitted to children j at age 25 as follows: $z_1^j = \rho_h z_{36}^i + \nu^j,\ \nu^j \sim N(0, \sigma_h^2)$, as parents are 35 years older than their children.

[2]In Japan, the statutory eligibility age for full public old-age pension benefits is 65 for both the National Pension (*Kokumin Nenkin*) and the Employees' Pension Insurance (*Kōsei Nenkin*). Eligibility for early claiming begins at age 60 with actuarially reduced benefits, while delayed claiming is permitted up to age 75 with actuarially increased benefits. Although many firms historically imposed mandatory retirement at age 60, current regulations require employers to offer continued employment opportunities up to age 65, creating a distinction between formal retirement practices and statutory pension eligibility. According to OECD (2025), the 2024 effective retirement ages for men and women were 66.9 and 67.4, respectively. Since claiming benefits and withdrawing from the labor market occur at the same age in our model, we take this age to be 67.

3.4 Preferences

Preferences are time separable, with a constant discount factor β. The period utility function from consumption is given by $U(c) = (c^{1-\gamma} - 1)/(1-\gamma)$.

People derive utility from holding onto assets because they turn into bequests upon death. This form of 'impure' bequest motive implies that an individual cares about total bequests left to their children, but not about the consumption of their children. The utility from bequests b is $\phi(b) = \phi_1[(b + \phi_2)^{1-\gamma} - 1]$. The term ϕ_1 measures the strength of bequest motives, while ϕ_2 reflects the extent to which bequests are luxury goods. If $\phi_2 > 0$, the marginal utility of small bequests is bounded, while the marginal utility of large bequests declines more slowly than the marginal utility of consumption. In the benchmark model, we set b as the bequest net of inheritance tax, b_n. We also consider the case in which gross bequests, b_g, enter the utility function. In that case, we set $b = b_g$.

3.5 The Household's Recursive Problem

We assume that children have full information about their parents' state variables and infer the size of the bequests that they are likely to receive based on this information. The potential set of a household's state variables is given by $x = (t, a, z, \tilde{y}, S_p)$, where t is household age (notice that in the presence of a fixed age gap, one's age is also informative about one's parents' age), a denotes the agent's financial assets carried from the previous period, z is the current earnings shock, and $\tilde{y}$ stands for the household's own annual accumulated earnings, up to a cap $\tilde{y}_c$ for counted earnings, which are used to calculate pension payments. The term S_p stands for parental state variables other than age and, more precisely, is given by $S_p = (a_p, z_p, \tilde{y}_p)$, where $\tilde{y}_p$ is the parent's accumulated earnings. It thus includes parental assets, current earnings, and accumulated earnings. When one's parent retires, z_p, or current parental earnings, becomes irrelevant and we set it to zero with no loss of generality.

From 25 to 66 years of age ($t = 1$ to $t = 42$), the agent works and survives for sure to the next period. Let $V_w(t, a, z, \tilde{y}, S_p)$ and $V_w^I(t, a, z, \tilde{y})$ denote the value functions of a working-age person whose parent is alive and dead, respectively, where I stands for 'inherited'. In the former case, the household's parent is still alive and might die with probability $1 - p_{t+35}$, in which case the value function for the orphan household applies, and assets are augmented by inheritances in per-capita terms. That is,

$$\begin{aligned} V_w(t, a, z, \tilde{y}, S_p) = \max_{c, a'} \Big\{ & U(c) + \beta p_{t+35} E\big[V_w(t+1, a', z', \tilde{y}', S_p')\big] \\ & + \beta(1 - p_{t+35}) E\big[V_w^I(t+1, a' + b_n/N, z', \tilde{y}')\big] \Big\}, \end{aligned} \tag{1}$$

subject to

$$c + a' = (1 - \tau_l)wy - \tau_s \min(wy, \tilde{y}_c) + Ra, \tag{2}$$

$$a' \geq 0, \tag{3}$$

$$\tilde{y}' = \big[(t-1)\tilde{y} + \min(wy, \tilde{y}_c)\big]/t, \tag{4}$$

$$\tilde{y}'_p = \begin{cases} \big[(t+34)\tilde{y}_p + \min(wy_p, \tilde{y}_c)\big]/(t+35) & \text{if } t < 43 \\ \tilde{y}_p & \text{otherwise} \end{cases} \tag{5}$$

$$b_n = b_n(S_p), \tag{6}$$

where N is the average number of children, b_n is bequests net of taxes, τ_l is the labor income tax rate, τ_s is the payroll tax (premium) for public pensions, and $\tilde{y}_c$ is the maximum taxable income for public pensions above which the payroll tax rate is zero.

In (2), $R = 1 + (1-\theta)(1-\tau_k)r_k + \theta(1-\tau_d)r_d$.[3] That is, a fraction θ of the assets must be held as government debt, d, with an exogenous return r_d and a tax on interest income τ_d, and the rest of the household's portfolio consists of titles to physical capital k, the return to capital r_k and the tax on the return to capital τ_k. Put differently, we assume that $d/a \geq \theta$.

The expected values of the value functions are taken with respect to (z', z'_p), conditional on (z, z_p). The agent's resources depend on labor endowment y and asset holdings a.

Average yearly earnings for children and parents evolve according to equations (4) and (5), respectively. Equation (6) gives the net bequest implied by the parent's optimal decision rule.

The value function of an agent who is still working but whose parent is dead is

$$V^I_w(t, a, z, \tilde{y}) = \max_{c,a'} \Big\{ U(c) + \beta E\big[V^I_w(t+1, a', z', \tilde{y}')\big] \Big\}, \tag{7}$$

subject to (2) and (3).

From 67 to 94 years of age ($t = 43$ to $t = 70$), the agent is retired and receives public pension benefits, and their parent is already deceased. They face a positive probability of dying, in which case they derive utility from bequeathing the remaining assets.

$$V_r(t, a, \tilde{y}) = \max_{c,a'} \Big\{ U(c) + \beta p_t V_r(t+1, a', \tilde{y}) + (1 - p_t)\phi(b) \Big\}, \tag{8}$$

subject to (3),

$$c + a' \;=\; Ra + (1 - \tau_l)P(\tilde{y}), \tag{9}$$

[3] Since we set $r_d = 0$, debt service is zero and τ_d does not affect the government budget; we retain it in the notation because the portfolio return is defined for general r_d. In the closed economy $D = 1.21\,Y$ rises with output. With zero growth, however, debt is constant in any stationary equilibrium, so there is no recurring revenue from issuance. The one-time proceeds from moving to a higher D arise only along the transition, which our steady-state comparisons do not include.

$$b_n = \begin{cases} a' & \text{if } a' < x_b, \\ (1-\tau_b)(a'-x_b)+x_b & \text{otherwise,} \end{cases} \tag{10}$$

and, in the case of net bequest motives,

$$b = b_n, \tag{11}$$

while in the case of gross bequest motives,

$$b = b_g = a', \tag{12}$$

regardless of the structure of the estate tax.

We focus on a stationary equilibrium concept in which factor prices and the age-wealth distribution are constant over time. Due to space constraints, the formal definition of a stationary equilibrium is in Appendix A.

3.6 Two Returns, and Two Polar Cases

Two features of the asset side of the model require comment, because they shape how the policy experiments should be read.

The portfolio return. Households hold two assets with different returns. Claims on physical capital earn the net marginal product $r_k = \alpha Y/K - \delta$, while a fraction θ of each household's assets must be held as government debt, which pays the exogenous rate r_d. The return that enters the household's Euler equation is therefore neither r_k nor r_d but the portfolio average

$$r_p \;=\; \underbrace{\frac{W-D}{W}}_{=\,1-\theta}\, r_k \;+\; \underbrace{\frac{D}{W}}_{=\,\theta}\, r_d, \tag{13}$$

where W is aggregate household net worth and D is government debt, so that $\theta = D/W$ is the debt share of the portfolio. The after-tax counterpart, which is what actually enters the household budget constraint (2) as $R-1$, is $(1-\theta)(1-\tau_k)\,r_k + \theta(1-\tau_d)\,r_d$.

The distinction matters for the experiments because fixing r_k does *not* fix r_p. When a reform raises W, the share of debt in household portfolios falls and, since $r_d < r_k$, the average return households earn rises endogenously, as (13) makes clear. In a model without government debt in household portfolios, $W = K$, $\theta = 0$ and $r_p \equiv r_k$, so the channel is absent; Section 6.2 shows that it accounts for about two-thirds of the response we report in our baseline case.

Two polar cases. The experiments below are reported under two polar assumptions about where additional household saving goes.

In the *small open economy*, which we take as the baseline, r_k is pinned at the world rate. The capital–labor ratio, the capital stock, output and the wage are then fixed, and household net worth satisfies

$$W \;=\; K \;+\; D \;+\; \text{NIIP},$$

where the net international investment position adjusts to absorb any difference. Additional saving induced by a reform accrues as net foreign assets rather than domestic capital. In the benchmark, Japan's existing net foreign assets are included in K, as in the capital series we use for calibration (Section 4), so NIIP is zero in the initial steady state and NIIP after a reform measures the *change* in the net foreign asset position.

In the *closed economy*, NIIP $\equiv 0$ and $W = K + D$, so the additional saving must be absorbed by domestic capital and r_k falls to clear the capital market. This is the standard case in the literature, and the one used by De Nardi and Yang (2016); it is a natural choice for the United States, which produces about a quarter of world output.

Neither polar case is literally correct. Japan is neither closed nor a price taker at an exogenous world rate for the entirety of its saving. We take the small open economy as the baseline, for the reasons below, but also report the closed economy, because the gap between them is large (Section 7) and the reader should be able to see how much turns on the choice.

Japan's size points toward the small open economy. Japan produces about 4 percent of world output. In our baseline, abolishing the tax raises Japan's net foreign assets by about 60 percent of its output, roughly 2.4 percent of world output; with a world capital–output ratio of about three and a capital share of one third, this lowers the world return to capital by only a few basis points. What the small open economy leaves out is home bias: if additional saving is not freely invested abroad, part of it must be absorbed at home, and the closed economy is the limiting case. The two cases also differ in the debt requirement. Because government debt is proportional to output, the closed economy, where output rises, also requires households to hold more zero-yield debt, which contributes to its smaller response.

4 Calibration

A model period corresponds to one year and we focus on steady states. We choose the parameters of our model in two stages. In the first stage, we pick parameters that are based on economic statistics from the data, as well as parameters, such as relative risk aversion, that are consistent with the literature. In the second stage, we jointly estimate the remaining six parameters by minimizing a loss function based on the distance between six model moments and their data counterparts. Table 1 lists the parameters we use in the first stage; the parameters calibrated in the second stage are given in Table 2.

4.1 First-Stage Calibration

Technology We abstract from technological growth and set capital's share of income, α, and the depreciation rate to 0.3783 and 0.0842, respectively, following Hansen and İmrohoroğlu (2023).

	Parameters	Value	Sources
Production			
α	capital income share	0.3783	Hansen and İmrohoroğlu (2023)
δ	depreciation	0.0842	Hansen and İmrohoroğlu (2023)
Government policy			
τ_k	capital income tax	0.2196	Capital tax revenue/GDP, 2019
τ_d	bond income tax	0.20	Statutory withholding rate
$P(\widetilde{y})$	public pension benefit	see text	See text
τ_s	public pension tax	18.3%	Statutory contribution rate
Demographics			
n	annual population growth	0.0%	Assumption
p_t	survival probability	see text	IPSS cohort life tables (1990 cohort)
Preferences			
γ	risk aversion coefficient	2.0	Standard in the literature
Labor earnings			
ϵ_t	age-efficiency profile	see text	Yamada (2025)
ρ_z	AR(1) persistence	0.87	Kitao, Suzuki, and Yamada (2025)
σ^2_μ	shock variance	0.0910	Kitao, Suzuki, and Yamada (2025)
Inheritance Transmission			
ρ_h	correlation	0.33	Ueda (2015)
σ^2_h	variance	0.37	De Nardi (2004)

Table 1: Exogenous parameters used in the benchmark model.

Government Our calibration of the Japanese government's spending and taxation follows Hansen and İmrohoroğlu (2023), extended through 2023.[4] The capital–output ratio is 2.72 and the ratio of net government debt to output is 1.21, so that household net worth is $W/Y = 3.93$ and the share of debt in household portfolios, $\theta = D/W$, is 0.308 in the data (0.307 in the model, where $W/Y = 3.94$).

The capital income tax rate, τ_k, is set to 21.96%, the rate that matches the ratio of capital income tax revenue to GDP in 2019. The tax on interest earnings from government bonds is 20%. The payroll tax to finance public pensions is taken as 18.3%. The labor income tax rate, τ_l, is endogenously determined to satisfy the government's budget.

[4] Hansen and İmrohoroğlu (2023) follow Hayashi and Prescott (2002) and treat government investment as part of government consumption, and hence as part of G. They also include the Japanese current account in their definition of investment, to account for the international asset holdings of Japanese households; their measure of output is therefore GNP rather than GDP. We follow these conventions in constructing the empirical calibration targets. In the model, $Y = F(K, L)$ denotes production, and empirical benchmark GNP provides the scale used to translate model quantities into yen. Unless otherwise stated, model ratios use Y, not national income including foreign earnings. In the small-open-economy inheritance-tax experiments, production remains at its benchmark level Y_0. The reported ratios therefore use a fixed benchmark-GNP scale, rather than reform-specific GNP: foreign assets generate additional resident income even though domestic production is unchanged.

Public pension benefits, $P(\tilde{y})$, are given by a lump-sum component (basic pension) and a part that depends on average past income:

$$P(\tilde{y}) = p_0 + p_1 \tilde{y}, \tag{14}$$

where $\tilde{y}$ is an individual's average past earnings. We take p_0 to be 15% of average earnings per person and calibrate the replacement rate p_1 so that the public pension system is balanced, given the exogenous payroll tax rate of $\tau_s = 0.183$.[5]

Demographics The population growth rate n is assumed to be zero. The p_t's are the vectors of conditional survival probabilities for people 67 and older, and are set to the survival probabilities for people born in 1990, using the cohort life tables from the Institute of Population and Social Security Research (IPSS), Ministry of Health, Labor and Welfare, Japan. The survival probabilities are plotted in Figure 1.[6]

The wealth–output target Because the discount factor is chosen to match household net worth relative to output, and because that ratio also determines the debt share θ and hence the wedge between r_k and r_p, it is worth setting out how we construct it.

The capital stock is built in the tradition of Hayashi and Prescott (2002), who accumulate Japanese investment flows and add the current account to GDP, so that the resulting output concept is GNP and the capital concept includes the claims Japanese residents hold on the rest of the world. Hansen and İmrohoroğlu (2023) update that series; we extend it through 2023 and obtain $K/Y = 2.72$. To this we add net government debt, which stood at 1.21 times output in 2023 and is held in household portfolios. Household net worth relative to output is therefore $W/Y = 3.93$. The benchmark model delivers $W/Y = 3.94$ and $K/Y = 2.73$. Because Japan's existing net foreign assets are already included in K, the benchmark net international investment position is zero, and only *changes* in the net foreign asset position are interpreted in the experiments below; Japan's actual position is on the order of 80% of GDP.

[5]Public pension contributions in both Japan and the United States are subject to earnings caps, though their relative generosity differs substantially. In Japan, Employees' Pension Insurance contributions are assessed on a capped standard monthly remuneration, with the maximum bracket currently set at approximately ¥650,000 per month (about ¥7.8 million annually), corresponding to roughly 1.6 times average income. As a result, pension contributions cease to rise at relatively moderate income levels. In contrast, the U.S. Social Security payroll tax applies up to a much higher annual taxable maximum (approximately $170,000), equivalent to about 2.4 times average individual earnings, and binds only for the top 6–7 percent of earners. This difference implies a broader contributory base and a less sharply declining effective contribution rate at higher incomes in the U.S. relative to Japan.

[6]Our targets are drawn from the years in which each is best measured: earnings inequality from the 2004 and 2014 waves, the fiscal and estate-tax aggregates from 2023, and the capital income tax rate from 2019. The survival schedule is the 1990 birth cohort's, which describes mortality at ages 67 and above several decades hence rather than that of current decedents. We use it because the steady state is meant to describe the post-2015 tax regime as a long-run position rather than any particular year, and we report the 2070 birth cohort's schedule in Section 10 as the alternative. Survival to 67 is certain in the model, so there are no accidental bequests from earlier deaths and every heir is at least 32 when they inherit; both omissions reduce the accidental component we measure and delay inheritances relative to the data.

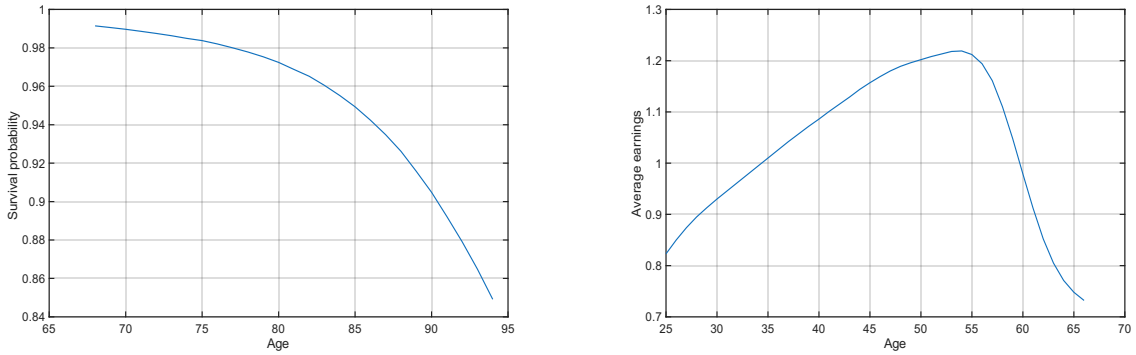
Over our sample period the return on Japanese government debt is close to zero, and we set $r_d = 0$ accordingly. Equation (13) then reduces to

$$r_p \;=\; r_k \frac{W - D}{W} \;=\; (1 - \theta)\, r_k, \tag{15}$$

which is the form we use throughout. At our benchmark calibration $\theta = D/W = 1.21/3.94 = 0.307$ and $r_k = 5.42\%$, so $r_p = 0.693 \times 5.42\% = 3.76\%$. The gap of roughly 1.7 percentage points between the two is the drag imposed by the zero-yielding debt, which accounts for about 31% of household portfolios.

Preferences We take the risk aversion coefficient, γ, to be 2.0. This parameter is the inverse of the intertemporal elasticity of substitution (IES) in consumption, whose estimates vary substantially: about 0.2 in Cashin and Unayama (2016) and Yogo (2004), 0.67 in Attanasio and Weber (1993), and 1.0 in Vissing-Jørgensen (2002). We take a central value of 0.5, which gives $\gamma = 2.0$.

Earnings Risk and Intergenerational Transmission of Ability The deterministic age-profile of labor earnings ϵ_t is taken from Yamada (2025) and is plotted in Figure 1. Since we impose mandatory retirement at the age of 67, we set $\epsilon_t = 0$ from that age on ($t \geq 43$).



(a) (b)

Figure 1: Panel (a): conditional survival probabilities by age, 1990 cohort. Panel (b): the age-efficiency profile of labor earnings.

To represent the earnings risk facing an individual, we assume an AR(1) process with persistence parameter 0.87 and shock variance 0.091, estimated by Kitao, Suzuki, and Yamada (2025) using prefecture-level tax return data.

We take the intergenerational correlation of earnings as 0.33 following Ueda (2015), and the innovation standard deviation as 0.6083, corresponding to a variance of approximately 0.37, following De Nardi (2004).[7]

[7]The latter is estimated for the United States; we are not aware of a comparable Japanese estimate of the dispersion of the intergenerational ability innovation.

To capture the very top of the earnings distribution, which household surveys measure poorly, we discretize the AR(1) process into four states using Tauchen (1986) and add a fifth, super-high state that can be reached only from the fourth. Its level and three transition probabilities are chosen to match the earnings Gini and the shares of total earnings received by the top 1, 5, 10, 20, 40 and 60 percent, computed over the full population including retirees. The calibration implies that the super-high state is transitory: a worker who reaches it stays for one year and then returns mostly to the lowest earnings state. The process reproduces the earnings Gini (0.562 against 0.552 in the data) and the top 1% and 5% earnings shares (6.0% and 19.9% against 6.1% and 19.5%). Appendix C.1 gives the earnings grid, the transition matrices, and the full comparison of earnings moments.

4.2 Second-Stage Parameters and Targeted Moments

Table 2 lists the six parameters that we calibrate in the second stage, together with the six data moments that they target, for the benchmark specification and for the two alternatives we use below: one in which parents value the pre-tax rather than the after-tax estate (gross bequests), and one in which the voluntary bequest motive is absent altogether. We choose β to match a wealth–output ratio of 3.93, constructed as described above. We calibrate the benchmark under the closed-economy condition $W = K + D$, choosing the discount factor β jointly with the other second-stage parameters to match the wealth–output target. We then set the small open economy's world return equal to the capital return in this reference equilibrium and retain the calibrated parameters. The reference allocation therefore also constitutes a small-open-economy benchmark with zero NIIP and identical factor prices. This construction aligns the two cases before reform; it does not match Japan's observed foreign position or domestic capital stock separately. With domestic capital and debt fixed, subsequent changes in household wealth accrue as changes in NIIP.

In the cases in which a bequest motive is present, we choose ϕ_1 and ϕ_2 to match the bequest–wealth ratio of 0.91% and the 90th percentile of the distribution of bequests across *all* decedents, which is 8.0 times average income, or about ¥40 million. Because roughly a tenth of decedents leave a taxable estate, this percentile sits close to the taxability threshold, which is why it is informative about ϕ_2.

Two percentiles need to be kept apart here, since both appear below. The targeted moment is the 90th percentile of bequests over all decedents, ¥40 million. The 90th percentile of *taxable* estates is a different and much larger object, approximately ¥250 million; Appendix B reports the underlying distribution from the 2023 National Tax Agency statistics.[8] We use the second figure only to locate the relevant bracket of the statutory rate schedule, as we do in the next subsection.

[8] *Source:* National Tax Agency (2025).

Targeted moment	Data	Benchmark (net bequests)	Gross bequests	No bequest motive
Wealth–output ratio, W/Y	3.93	3.94	3.94	3.93
Bequest flow / wealth, B/W	0.91%	1.00%	0.98%	0.67%
P90 of bequests / average income	8.00	6.63	6.80	5.23
Fraction of decedents taxed	9.90%	9.38%	9.75%	10.80%
Inheritance tax revenue / Y	0.50%	0.50%	0.50%	0.50%
Government spending / Y	22.0%	22.0%	22.0%	22.0%
Implied parameters				
β, discount factor		0.9690	0.9702	0.9839
ϕ_1, strength of bequest motive		−43.4035	−43.5393	—
ϕ_2, bequest utility shifter		¥12M	¥13M	—
τ_b, inheritance tax rate		39.5%	45.9%	62.1%
x_b, effective exemption		¥34.8M	¥36.2M	¥25.0M
τ_l, labor income tax rate		24.8%	24.8%	24.8%

Table 2: Second-stage calibration: targeted moments and calibrated parameters.

Notes: The top panel reports the six targeted data moments and their model counterparts for the benchmark model, in which parents value the after-tax estate (net bequests); for the alternative in which they value the pre-tax estate (gross bequests); and for the model with no voluntary bequest motive. The 90th percentile of the bequest distribution is expressed as a multiple of average income. The bequest–wealth ratio and the 90th percentile are both governed by ϕ_1 and ϕ_2 and pull against each other: matching the level of the flow calls for a stronger motive near the median, while matching its upper tail calls for a larger luxury threshold. The estimation trades them off, overshooting the flow by 10% and undershooting the percentile by 17%. Fitting both exactly would require a third parameter in the bequest function. The bottom panel reports the corresponding calibrated parameters; the exemption level x_b and the bequest utility shifter ϕ_2 are in millions of yen. All three specifications are calibrated to the same two administrative facts about the inheritance tax, so that the comparison between the first and third columns is the decomposition of Hendricks (2002).

The discount factor affects savings and average wealth in the economy. The term ϕ_1 measures the strength of bequest motives, and we therefore discipline it with the flow of aggregate bequests relative to wealth. The term ϕ_2 reflects the extent to which bequests are luxury goods, affecting the upper tail of the bequest distribution. The calibrated ϕ_2 of ¥12 million—roughly 2.4 times average income—implies that the marginal utility of bequests is bounded for small estates, so that a large fraction of households optimally leave bequests of modest size, while the desire to bequeath declines only slowly with wealth for the richest households. The calibrated ϕ_1 is −43.40. The negative sign is a consequence of $\gamma = 2$: since $1-\gamma = -1$, the term $(b+\phi_2)^{1-\gamma}$ is decreasing in the bequest, so a negative ϕ_1 is what makes the bequest component of utility increasing and concave in b, in exactly the way that $c^{1-\gamma}/(1-\gamma)$ is negative but increasing in consumption at this value of γ.

The estimated bequest function implies that a household in its last year of life leaves nothing if its resources fall below about a third of average income (¥1.8 million) and devotes 87 of every 100 additional yen above that to its estate; Appendix C.2 derives these statistics and the bunching of estates at the exemption.

Although many experts agree that effective inheritance taxation can be reduced substantially by appropriate estate planning and valuation—in Japan, most notably through the preferential valuation of real estate—there is considerable uncertainty about how much households can and do reduce the inheritance tax burden through both legal and illegal means. There is, in contrast, no dispute about the observed revenue from the inheritance tax and the fraction of decedents whose estates pay it. We therefore choose the tax parameters τ_b and x_b to match the ratio of inheritance tax revenue to output (0.50%) and the fraction of decedents whose estates pay inheritance tax (9.9%), both from the National Tax Agency.[9] The implied effective exemption level turns out to be ¥34.8 million, which is close to the statutory basic exemption of ¥30 million plus ¥6 million per statutory heir (¥36–42 million for a decedent with one or two heirs; see Appendix B). The implied effective flat marginal rate is 39.5%, which lies within the statutory progressive schedule that rises from 10% to 55%.[10] Neither of these features of the tax code is targeted. Given that our model matches asset holdings well, and given the considerable uncertainty about effective avoidance and evasion, we see this as a useful way to proceed.

We choose τ_l to balance the government budget constraint, given a ratio of government spending to output of 22%; the resulting labor income tax rate is 24.8%.

In the model without voluntary bequest motives, we choose β, τ_b, x_b and τ_l to match the wealth–output ratio, the ratio of inheritance tax revenue to output, the fraction of estates that pay inheritance taxes, and the ratio of government spending to output. This recalibration is deliberate: the comparison asks how much bequest flow an economy without a voluntary motive generates while remaining disciplined by the same wealth and inheritance-tax targets. Following Hendricks (2002), we use this comparison as a model-based decomposition. Because removing the motive also requires adjustments to the discount factor and effective tax parameters, the difference in flows is not a fixed-parameter estimate of the voluntary motive's contribution. Revenue relative to output

[9]The National Tax Agency's 9.9% is the ratio of decedents filing returns *on which tax is assessed* (155,740 in 2023) to all deaths (1,576,016). Returns filed with no tax due—38,121 in 2023, principally estates reduced to zero by the spousal credit—are reported separately and excluded; including them would give 12.3%. Revenue is likewise measured on returns with tax assessed, so our two targets refer to the same population of estates. Since our model contains no spousal credit, every estate above x_b pays, and 9.9% is the appropriate target.

[10]Two features of the Japanese tax lie outside our single-person household structure. Most first deaths pass to a surviving spouse, who can deduct the larger of the statutory share or ¥160 million, so much of a married couple's wealth is taxed only at the second death; our model therefore dates transfers earlier in the life cycle than the data. And liability is computed on a notional division among statutory heirs, about 2.4 on average, rather than on the whole estate, whereas in our model τ_b applies to the whole estate above x_b. This is why we compare τ_b only with the range of statutory rates, not bracket by bracket; Appendix B describes the computation.

is matched exactly in both at 0.50%; the fraction of decedents taxed is hit less precisely, 9.38% in the benchmark and 10.80% without a voluntary motive against a target of 9.90%.

An inspection of Table 2 shows that the model without a voluntary motive cannot match the flow of aggregate bequests relative to wealth: it generates a bequest–wealth ratio of 0.67%, about three-quarters of the 0.91% observed in the data and two-thirds of the 1.00% the benchmark generates, and it understates the 90th percentile of the bequest distribution (5.23 versus 8.0 times average income). As a consequence, it overstates the inheritance tax rate needed to raise the observed revenue, setting it at 62.1%—above the top statutory marginal rate of 55%, which applies only to taxable shares above ¥600 million. The higher rate arises mechanically: the flow of bequests is too small, yet revenue must still match the data. Matching the same wealth–output ratio without a bequest motive also requires a substantially higher discount factor, $\beta = 0.9839$ against 0.9690 in the benchmark, since the warm glow is no longer available to generate saving late in life. The benchmark model, in contrast, matches four of the six targets closely and the remaining two—the bequest–wealth ratio, 1.00% against 0.91%, and the 90th percentile of bequests, 6.63 against 8.00 times average income—to within ten and seventeen percent, respectively, with an effective rate and exemption level that line up with the statutory tax code.

The gross-bequest specification, in which parents value the estate before tax rather than after, matches the targets about equally well but requires a higher effective rate, 45.9%, and a slightly higher exemption. Its preference parameters are almost identical to the benchmark's—$\beta = 0.9702$ against 0.9690, $\phi_1 = -43.54$ against -43.40, $\phi_2 = ¥13$ million against ¥12 million—so the difference between the two specifications in what follows comes from how the tax enters the bequest utility function rather than from any material difference in the strength or shape of the bequest motive. The reason is that when the tax does not enter the bequest utility function, a given rate does less to induce parents to save in order to protect the after-tax estate, so a higher rate is needed to raise the same revenue from a similar flow. We report results for both specifications throughout, since there is no decisive evidence on which is the better description of Japanese bequest behavior.

5 Numerical Results

This section highlights the implications of the calibrated model, none of which are targeted in the calibration; they therefore serve as a test of the model.

5.1 The Wealth Distribution in the Cross Section and Over the Life Cycle

Table 3 reports the wealth distribution. The first line refers to the Japanese data of Kitao and Yamada (2025) and shows that wealth in Japan is unevenly distributed, though

considerably less so than in the United States.[11] The wealthiest 1% of Japanese households hold 9.8% of net worth (compared with roughly 35% in the U.S.), the wealthiest 5% hold 28.7%, the wealthiest 10% hold 43.8%, and the wealthiest 20% hold 63.7%, with an overall wealth Gini of 0.668.

		Share of wealth held by the top (%)					
	Gini	1%	5%	10%	20%	40%	60%
Data: Kitao and Yamada (2025)	0.668	9.8	28.7	43.8	63.7	85.2	95.5
Benchmark (net bequests)	0.61	17.4	29.3	41.7	60.3	83.8	95.3
Gross bequests	0.60	15.9	27.9	40.4	59.4	83.4	95.2
No voluntary bequests	0.55	12.8	24.9	37.0	55.4	80.1	93.6
Difference: benchmark less no-motive	+0.06	+4.6	+4.4	+4.7	+4.9	+3.7	+1.7

Table 3: Percentage of total wealth held by households in the top percentiles.

Notes: None of these moments are targeted in the calibration. First line: Japanese data (average of the 2004 and 2014 values) from Kitao and Yamada (2025). The three model rows are the benchmark, in which net-of-tax bequests enter the utility function; the gross-bequest specification; and the model without a voluntary bequest motive. The last line is the difference between the benchmark and the no-voluntary-bequest economy, both of which are disciplined by the same inheritance-tax targets. Households aged 25 are excluded from the model calculations because they are born with zero net worth.

The second line reports the corresponding numbers for the benchmark model with intergenerational links and voluntary bequest motives, with bequests net of taxes entering the utility function. The share of the top 5% is 29.3% in the model versus 28.7% in the data, the top 10% share is 41.7% versus 43.8%, the top 20% share is 60.3% versus 63.7%, the top 40% share is 83.8% versus 85.2%, and the top 60% share is 95.3% versus 95.5%.

The 0.61 model Gini is somewhat below the data value of 0.668, and the model overstates the share of the very richest 1% (17.4% versus 9.8%). The latter reflects the parsimony of our five-state earnings process, in which the entire upper tail of earnings is concentrated in a single super-high state. That state exists in order to generate estates large enough to reach the statutory top brackets, and the two administrative moments we target—revenue relative to output and the fraction of decedents taxed—discipline the upper tail of bequests directly. Both features concentrate the action at the very top, which is why the model matches the top 5% share almost exactly and the top 10, 20 and 40 percent shares to within one and a half to three and a half percentage points, while overshooting the top 1 percent. We view the fit of the upper half of the distribution as the more informative test.

The fourth line reports the wealth distribution generated by the model without voluntary bequests, recalibrated as discussed previously. This version generates wealth holdings

[11] Since all of our 25-year-olds are born with zero net worth, we exclude them from our calculations of wealth inequality.

in the hands of the richest 1% (12.8%) that exceed the share of earnings of the richest 1% (6.1%). The key mechanism is that the earnings super-rich face a substantial probability of sliding into a much lower earnings state, and thus save at very high rates to smooth consumption over time. However, despite this additional saving motive for high earners, the model without a voluntary bequest motive falls short of the concentration generated by the benchmark everywhere above the median: its Gini is 0.55 (versus 0.61 in the benchmark), and its top 1%, 5%, 10% and 20% shares are 12.8%, 24.9%, 37.0% and 55.4%, each four to five percentage points below the benchmark. The gap closes as one moves down the distribution—1.7 percentage points at the top 60%—which is the sense in which the voluntary motive does its work above the median. The gross-bequest specification sits between the two, closer to the benchmark.

The comparison between the two models highlights the role of the voluntary bequest motive, calibrated as a luxury good, in generating a higher concentration of wealth in the hands of the richest few and raising overall wealth inequality. The intuition is that this kind of bequest motive raises the saving rate of the rich, who thus leave larger estates to their children; the children in turn also save more, thus increasing wealth concentration. This mechanism, documented for the United States by De Nardi and Yang (2016) and for Japan by Ono (2024), is also clearly at work in our Japanese calibration.

It is worth noting how our estimate of the magnitude compares with Ono (2024), who asks the positive version of this question for Japan in the same modeling lineage but with an independent calibration and an earlier vintage of the Kitao–Yamada data. He finds that intergenerational links raise the wealth Gini by 0.028 and the top 1% share by 0.6 percentage points; we find 0.06 and 4.6 points. The sign and the pattern agree—concentration rises everywhere above the median—which is consistent with the mechanism, though the magnitudes differ considerably: our motive does about eight times as much to the top 1% share and about four times as much to the top 5%, while at the top 20% the two estimates are within a factor of three. The two exercises are not constructed to be comparable in scope, so the agreement in sign is as much as the comparison supports. Our fit errors run in the opposite direction from his: he undershoots the top 1% share (6.2 against 10.2 in his data) while we overshoot it, and he is six to eight percentage points short over the top 5 to 20 percent, where we are close. The difference reflects the two features described above, both of which are consequences of our targeting the administrative moments of the inheritance tax, which he has no reason to target.

Figures 2 and 3 compare the model's life-cycle profiles of average net wealth and of within-age-group wealth inequality with their empirical counterparts; neither is targeted in the calibration. The model reproduces the hump-shaped profile of average net wealth, tracking the data closely over the accumulation phase and peaking at about the same ages, 60 to 70, although at a somewhat higher level. After age 75 the model decumulates faster than in the data, where average net wealth remains remarkably flat into the mid-80s; this is a well-known challenge for life-cycle models, and the comparison of the two model economies shows the bequest motive working in the right direction, as the benchmark decumulates visibly more slowly after age 70 than the economy without voluntary

bequests.

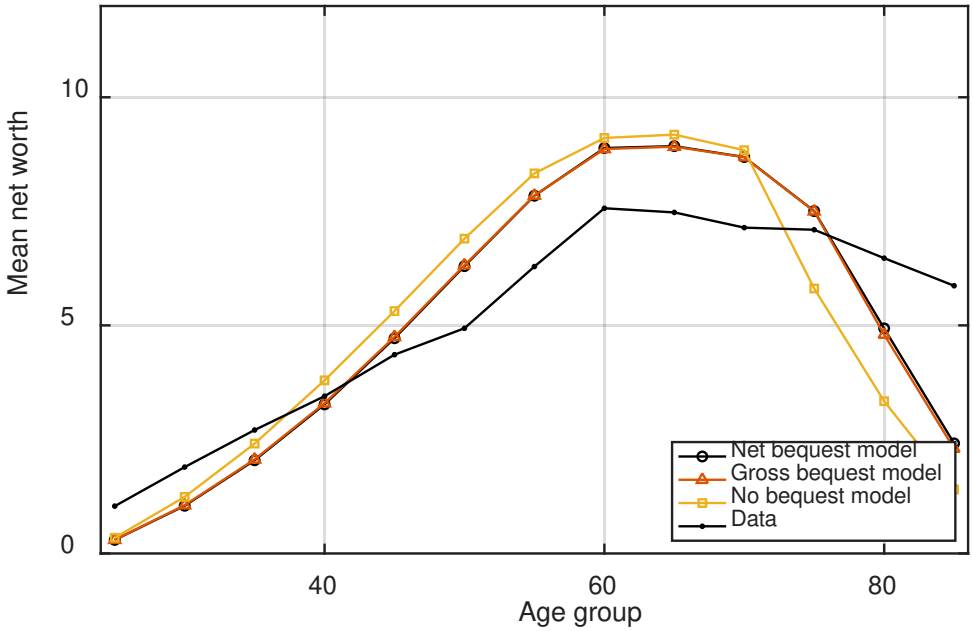


Figure 2: Average net wealth over the life cycle: model and data.

Notes: Data are the average of the 2004 and 2014 values from Kitao and Yamada (2025). Households aged 25 are excluded because they are born with zero net worth. This moment is not targeted in the calibration.

The model also captures the level and shape of within-age inequality at both ends of the life cycle: the wealth Gini among the young is about 0.72 in both model and data, and the model reproduces the rise in inequality at advanced ages, reaching 0.60 at ages 85 and older against 0.63 in the data—a rise that is largely absent without the bequest motive.

Around retirement the model generates more equalization than the data, in which within-age inequality stays near 0.6 throughout middle age and rises at older ages; the gap plausibly reflects sources of wealth dispersion outside the model that are particularly relevant for Japan, notably heterogeneity in housing wealth—across owners and renters and across regions with very different land values—and the persistent earnings and pension gap between regular and non-regular employees, which a parsimonious five-state earnings process cannot fully capture. Overall, we read the two figures as supportive of the model's central mechanism: the voluntary bequest motive is what slows old-age decumulation and regenerates wealth dispersion among the elderly, and both margins move the model toward the data.

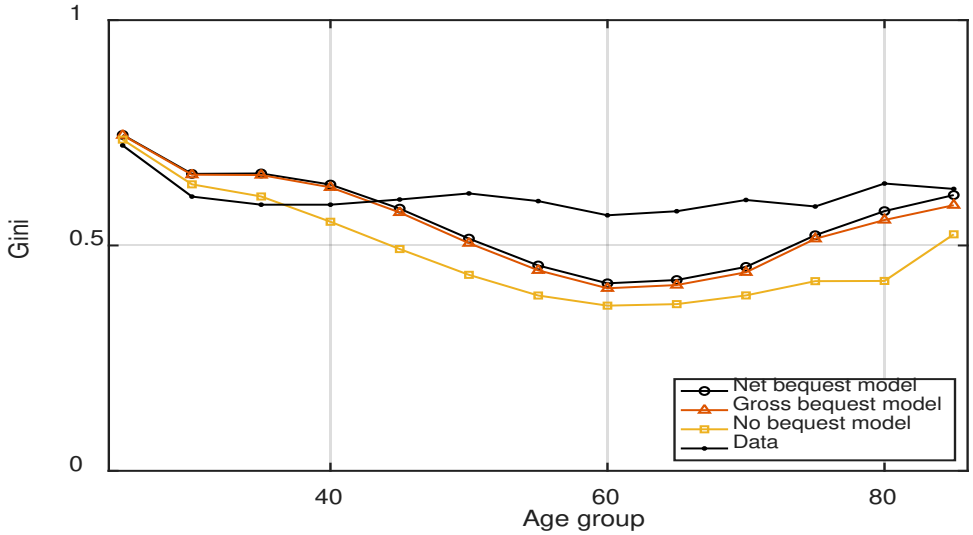


Figure 3: Wealth inequality within age groups over the life cycle: model and data.

Notes: Gini coefficient of net wealth within each five-year age group. Data are the average of the 2004 and 2014 values from Kitao and Yamada (2025). This moment is not targeted in the calibration.

The economy without a voluntary motive used in these comparisons generates about two-thirds of the benchmark bequest flow relative to wealth; Section 9.1 uses this comparison to decompose the bequest flow into accidental and intentional components.

5.2 The Importance of Parental Background

Parental background affects one's prospects in life through two channels. First, since richer parents leave larger bequests, it influences the amount of expected bequests that one will receive. Second, since one's initial earnings draw is correlated with one's parental earnings and is then persistent over time, it also influences one's lifetime earnings.

In this subsection, we discuss the value of being born to a family with different parental backgrounds (or earnings). Later, when evaluating various policy reforms, we assess to what extent inheritance, capital income, and labor income taxes can affect the luck (or lack thereof) of being born with a certain parental background rather than another one in our framework.

In our calibration, the earnings of both parents and children can assume five values, ordered from the lowest (1st) to the super-high (5th) earnings state. We perform our calculations of the value of being born in a certain parental socio-economic class as follows. Take a new worker with parental background i. Find the median of the other parental state variables (assets and associated lifetime earnings) with current parental earnings or background at the time when the child enters the labor market. Also find the median of the new workers' state variables (initial earnings) conditional on their parental earnings or background. Take the corresponding value function for all of these state variables. Repeat this process for a new worker with parental background j. Compare the two

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Panel A: bequests and human capital inheritance				
2nd	0.04	—	—	—
3rd	1.98	1.85	—	—
4th	2.33	2.21	0.27	—
5th	18.03	17.91	15.52	14.88
Panel B: bequests only				
2nd	0.04	—	—	—
3rd	0.37	0.27	—	—
4th	0.68	0.57	0.27	—
5th	13.94	13.80	13.35	12.92

Table 4: The value of parental background in the benchmark economy.

Notes: Each entry is the one-time asset compensation, expressed as a multiple of average income, required to make a newborn worker born to the parental background in the row indifferent to being born to the background in the column. Panel A switches both the child's initial productivity and the expected bequest; Panel B switches the expected bequest while holding the child's initial productivity fixed. Parental earnings states are ordered from the lowest (1st) to the super-high (5th). When the child enters, parents are aged 60, and their shares in the five earnings states are 16.3%, 33.7%, 33.7%, 16.3% and 0.002%, so the fifth row describes a group of negligible size; the compensations in that row should be read as a property of the model's upper tail rather than as a statement about an identifiable population.

value functions, and compute the one-time asset compensation required to make the newborn worker born to a given family background indifferent to being born to a family with another level of parental earnings or background and divide by average income to normalize. One way of interpreting this comparison is that it calculates the value of being born to a typical background, conditional on parental socio-economic status, and all of the median associated state variables that go along with it, with a different parental socio-economic status and all associated other median state variables.

In Panel A of Table 4 we switch both initial human capital and expected bequests among the offspring of families from different parental backgrounds, thus evaluating the importance of bequests and of initial productivity together. In Panel B we switch only expected bequests, thus isolating the importance of wealth transmission across generations for a given level of initial productivity.

Panel A of Table 4 shows that newborn workers whose parents are in the super-high (5th) earnings state at age 60 need to be compensated by 18.03 times average income to be moved to the state of being born to a parent in the lowest earnings state, and by 17.91, 15.52 and 14.88 times average income to be moved to the 2nd, 3rd and 4th backgrounds, respectively. Outside of the super-high state, the advantages of a better background are far more modest: a newborn worker with parents in the 4th (high) earnings state requires

a compensation of only 2.33 times average income to be moved to the lowest background, and one with parents in the 3rd state 1.98 times average income. Within the bottom four backgrounds the differences are smaller still—moving from the 4th to the 3rd background is worth 0.27 times average income, and from the 2nd to the 1st only 0.04. The bottom four backgrounds are thus close to one another in terms of the lifetime utility they deliver to children, compared with the advantage of being born to super-high earning parents.

Panel B shows that, for those born to a super-rich family, expected bequests are the dominant source of this advantage: the bequests-only compensation for moving from the 5th to the 1st background is 13.94 times average income, or 77% of the total of 18.03, with the remaining 23% attributable to the inheritance of human capital. The pattern across the columns of Panel B is also informative. The bequests-only compensations for a 5th-state child are nearly the same whichever of the bottom four backgrounds they are moved to—13.94, 13.80, 13.35 and 12.92—because none of those four backgrounds leaves an estate large enough to matter. In Panel A, by contrast, the compensations fall visibly as the destination background improves, from 18.03 to 14.88. Since the bequests-only compensations fall only from 13.94 to 12.92, about two-thirds of that decline is the human capital channel.

This decomposition is quite different from the one that De Nardi and Yang (2016) obtain for the United States, where bequests account for roughly 40% of the advantage of a super-rich background and human capital inheritance for 60%. Two features of the Japanese calibration drive the difference: the intergenerational correlation of earnings is lower (0.33), so the super-high earnings state is transmitted to children less reliably, and the estates left by super-high earners are extremely large relative to average income. The importance of bequests also shrinks quickly as we move down the parental background distribution: for children of 4th-state parents, bequests account for 0.68 of the total compensation of 2.33 (29%), and for children of 3rd-state parents for 0.37 of 1.98 (19%), with human capital inheritance accounting for the rest. The two panels coincide in the 2nd state, where both channels are negligible, and also for the move from the 4th background to the 3rd, where the human capital difference between two adjacent middle states is essentially nil. The implication for policy is that the inheritance tax in Japan operates on the dominant channel of dynastic advantage, whereas in the United States it operates on the secondary one.

These calculations thus suggest that the value of being born into a family with a super-high socio-economic background is very large, and that parental background—and especially expected bequests—is an important determinant of the lifetime utility of the richest households in our calibrated model. In contrast, the effects of parental background are much smaller for households born to the low- and middle-income parental levels, with the importance of bequests becoming much smaller as the level of parental background decreases. Since in our economy the fraction of people born to super-high earning parents is negligible, for the vast majority of people the effect of parental background on lifetime expected utility, while non-trivial, is far smaller.

These properties of the benchmark economy matter for the policy experiments that

follow. Because the dynastic advantage documented above is financed almost entirely by the bequests of a very small group of households whose estates are very large, the exemption and the rate are not interchangeable instruments. A higher exemption at an unchanged rate relieves estates in near-inverse proportion to their size, so it leaves the tax's bite on the largest estates nearly intact while removing it from the broad base of moderately sized ones; a rate cut relieves every taxable estate in the same proportion. Sections 6 and 7 show that this distinction, and not the size of the aggregate response, is what separates the two reforms.

6 Reforming Inheritance Taxation in a Small Open Economy

We study two key margins of inheritance taxation: the threshold above which estates start being taxed, and the marginal tax rate that applies above that threshold. Modifying the exemption affects both the size of the estates that are hit by the tax and the burden on the estates that remain taxable. For example, reducing the exemption implies that smaller estates start being taxed, but also that previously taxed estates pay more, because their exemption is smaller. In contrast, changing the tax rate for a given exemption simply increases or decreases the burden on estates of the same size.

Changing inheritance taxation also has an effect on the distribution of estates, for two reasons. First, if those leaving estates care about the estate net of taxes, they will change their saving behavior and desired bequest when the tax changes. This effect is absent when gross bequests enter the utility function. Second, heirs receive different amounts net of taxes, which affects their own saving behavior and desired bequests, because a realistically calibrated bequest motive generates a non-homotheticity of saving in income and wealth.[12]

We consider three reforms. In each case the government budget is rebalanced by adjusting the labor income tax; Section 6.3 compares this with the alternative of adjusting the capital income tax, and the complete results for that case are in Appendix D.3.

The first reform eliminates inheritance taxation altogether, setting $\tau_b = 0$. This is the natural benchmark experiment, and the one studied for the United States by Cagetti and De Nardi (2009) and De Nardi and Yang (2016).

[12]The luxury bequest motive adds a distribution-sensitive saving channel to those already generated by borrowing constraints, uninsurable earnings risk, and differences in remaining lifetimes. Because the utility from bequests features a large shifter ϕ_2, households of modest means place little weight on bequests and consume most of an inheritance, while wealthy households save a large fraction of it toward bequests of their own. The marginal propensity to save therefore rises with wealth and permanent income. As a result, the inheritance tax affects capital accumulation not only by shrinking the total flow of transfers but also by reshaping their distribution across recipients with very different saving propensities—and, since today's net bequests influence tomorrow's desired bequests, the effect propagates across generations of a dynasty.

The second reform keeps the exemption unchanged at ¥34.8 million and halves the effective rate τ_b, from 39.5% to 20%.

The third reform holds the tax rate at its benchmark value but raises the exemption to ¥575 million, approximately 115 times average income—the level implied by the *current* U.S. federal estate tax exemption relative to U.S. average income. This experiment asks what would happen if Japan adopted a U.S.-style, narrow-base inheritance tax *as that tax exists today*. Given the empirical distribution of taxable estates (Appendix B, Table A2), an exemption of this size removes the great majority of currently taxable estates from the tax net, leaving roughly one taxable estate in forty—about two tenths of one percent of all decedents—subject to the tax, against 9.4% in the benchmark. It is worth being explicit that this is a far more radical narrowing of the base than anything in De Nardi and Yang (2016), whose policy experiment moves in the opposite direction, toward a broader base and a higher rate than their benchmark.

This section reports these reforms in the small open economy, which we take as our baseline: r_k is pinned at the world rate, so the capital–labor ratio, the capital stock, output and the wage are fixed, and the difference between household net worth and the sum of domestic capital and government debt accrues as net foreign assets. Section 7 reports the same reforms in the closed economy, which is also the basis for our comparison with the United States.

6.1 Aggregate Effects

Table 5 reports the aggregate effects. Because K and Y are fixed by construction, the aggregate quantity of interest is household net worth, and the counterpart of any increase in it is an increase in the net international investment position.

	τ_l	W/Y	$\Delta\text{NIIP}/Y$ (%)	B/W (%)	r_p (%)	T_b/Y (%)
Benchmark ($\tau_b = 39.5\%$)	0.248	3.940	—	1.003	3.759	0.499
Eliminate ($\tau_b = 0$)	0.245	4.542 (+15.3%)	+60.2	1.087	3.979	0.000
Cut the rate ($\tau_b = 20\%$)	0.247	4.211 (+6.9%)	+27.1	1.044	3.866	0.285
Raise exemption (¥575M)	0.249	4.172 (+5.9%)	+23.2	1.024	3.851	0.171

Table 5: Aggregate effects of inheritance tax reform in the small open economy, adjusting the labor income tax.

Notes: τ_l is the labor income tax rate that balances the government budget. W is aggregate household net worth, B the annual flow of bequests, and r_p the portfolio return defined in (13). $r_k = 5.42\%$ and $K/Y = 2.733$ are fixed in all rows, and so is Y, so every output-normalized entry uses the benchmark output as its denominator. $\Delta\text{NIIP}/Y$ measures the change in foreign assets from the zero-NIIP reference equilibrium. Foreign income increases by $r_k\Delta\text{NIIP}$, but the reporting denominator is not adjusted for that income.

Eliminating the inheritance tax ($\tau_b = 0$): Eliminating the inheritance tax raises household net worth by 15.3%, from 3.940 to 4.542 times benchmark output. Domestic capital and government debt are unchanged, so net foreign assets increase by 0.602 times benchmark output. On the benchmark-GNP scale used to translate model quantities into yen, this amounts to about ¥400 trillion, or $2.67 trillion at $1 = ¥150. These magnitudes are measured against benchmark GNP, not reform-specific national income.

The mechanism is the one emphasized by De Nardi and Yang (2016): removing the tax raises the return to leaving a bequest for households rich enough to have an active bequest motive, and their additional saving dominates in the aggregate. Two features of the Japanese setting amplify it. First, with r_k fixed there is no general equilibrium damping: the extra saving does not depress the return to capital. Second, and less obviously, the portfolio return *rises*. As aggregate net worth grows, the zero-yielding government debt becomes a smaller share of household portfolios, so by (13) the average return households earn increases, from 3.76% to 3.98%. A higher return to saving induces still more accumulation, so the two forces compound rather than offset. We return to the quantitative importance of this channel below.

The labor income tax *falls* slightly, from 24.8% to 24.5%, even though the government gives up 0.5% of output in inheritance tax revenue. The reason is that the larger wealth stock expands the base of the capital income tax by more than the revenue forgone.[13] The flow of bequests relative to wealth also rises, from 1.00% to 1.09%, because with r_k fixed nothing offsets the additional accumulation.

Lowering the inheritance tax rate to 20%: Halving the effective inheritance tax rate to 20% raises net worth by 6.9%, with a corresponding increase of 0.27 times output in the net foreign asset position. This is a little under half the effect of full elimination, which is roughly what the halving of the effective rate would suggest.

Raising the exemption to the current U.S. level: Raising the exemption to the current U.S. level, ¥575 million, while holding the rate at 39.5%, raises net worth by 5.9%, slightly less than the rate cut. The reform removes the tax from the margin for almost all households that leave taxable estates in the benchmark—the roughly 9% of decedents with estates between ¥34.8 and ¥575 million—while a 39.5% marginal rate continues to apply to the very largest. Because those very large estates are left by households with the strongest bequest motives and the highest saving rates, retaining the tax at the top costs a non-trivial part of the aggregate response. The same feature, as we show below, is what makes this reform distinctive distributionally.

[13] The arithmetic is worth setting out, because it is what makes the reform self-financing. The reform raises the net foreign position by 0.60 times output; at $r_k = 5.42\%$ that is 3.25% of output in additional income, which taxed at $\tau_k = 22\%$ yields 0.715% of output against the 0.50% of inheritance tax revenue forgone. The result therefore assumes that the marginal unit of foreign assets earns the domestic marginal product of capital and is taxed in full on a residence basis. Japan's realized return on its net foreign position is below its domestic return on capital, and foreign-source income is not taxed at the full domestic rate, so both assumptions push in the same direction and the fall in τ_l should be read as an upper bound.

6.2 The Portfolio Channel

It is worth isolating how much of the baseline response comes from the endogenous portfolio return rather than from the absence of general equilibrium damping, since the first channel is specific to an economy with a large stock of government debt in household portfolios.

We do this by computing the reforms in a pure partial equilibrium in which r_p is frozen at its benchmark value along with r_k, K, Y and w, so that the debt share θ is held fixed rather than being allowed to fall as wealth accumulates. Eliminating the tax then raises household net worth by 5.7%, against 15.3% when the portfolio return is allowed to respond. The portfolio channel therefore accounts for roughly two-thirds of the baseline response.

This is a feature of the Japanese environment that has no counterpart in the U.S. studies. In a model without government debt in household portfolios, as in De Nardi and Yang (2016), aggregate household assets equal the capital stock and $r_p \equiv r_k$, so fixing the return to capital fixes the return households earn. In Japan, where net government debt is 121% of output and makes up 31% of household portfolios, the two differ by 1.7 percentage points in the benchmark, and the gap narrows endogenously whenever a reform raises aggregate wealth.

The size of this channel depends on how household debt exposure scales with wealth, and our proportional requirement settles that by assumption. A fixed aggregate debt stock does not by itself imply that every household's marginal return rises when wealth increases; under a rule fixing each household's debt holding in level rather than in proportion, the marginal yen would earn r_k and would not vary with aggregate wealth, so no dilution would occur. The requirement is a modeling device rather than a measured feature of Japanese portfolios: most government debt is held by the Bank of Japan, banks and insurers rather than directly by households, whose exposure is largely indirect.

6.3 The Role of the Financing Instrument

The ranking of the three reforms depends on which tax instrument absorbs the change in inheritance tax revenue. Table 6 places the aggregate effects under the two financing schemes side by side; the full set of results when the capital income tax adjusts, including distributional effects and the value of parental background, is in Appendix D.3.

All six responses are positive, but their magnitudes and their ordering change completely. The source of the difference is a timing asymmetry between the two instruments. The inheritance tax is levied once per lifetime, on a stock that has already been accumulated; the capital income tax is levied every year, on the return to accumulating it. Replacing revenue from the first with revenue from the second therefore requires a large increase in the rate, and the increase falls directly on the saving margin the reform was meant to encourage.

Full elimination illustrates this most sharply. It gives up all 0.5% of output in inher-

	Labor income tax adjusts		Capital income tax adjusts	
	τ_l	ΔW	τ_k	ΔW
Benchmark	0.248	—	0.220	—
Eliminate ($\tau_b = 0$)	0.245	+15.3%	0.249	+1.2%
Cut the rate ($\tau_b = 20\%$)	0.247	+6.9%	0.234	+0.6%
Raise exemption (¥575M)	0.249	+5.9%	0.216	+8.1%

Table 6: Aggregate effects of the reforms under alternative financing instruments, small open economy.

Notes: ΔW is the percentage change in aggregate household net worth relative to the benchmark. In the left panel the labor income tax τ_l adjusts to balance the government budget; in the right panel the capital income tax τ_k adjusts. Complete results appear in Table 5 and, for the capital income tax, in Appendix Table A9.

itance tax revenue, and replacing that with capital income taxation requires raising τ_k from 22.0% to 24.9%. The higher annual tax on the return to saving nearly cancels the reform: net worth rises by 1.2% rather than 15.3%.

The exemption reform is the exception, and it is the exception for an instructive reason. It induces enough additional accumulation that the capital income tax base expands by more than the revenue forgone, so τ_k can *fall*, from 22.0% to 21.6%. The lower tax on saving then reinforces the direct effect of the reform, and the response grows from 5.9% to 8.1%. Under capital-tax financing, the exemption reform is therefore the largest of the three, while under labor-tax financing it is the smallest. The ordering of the reforms by aggregate effect is a statement about how the revenue is replaced, not about the strength of the bequest motive.

This margin receives little attention in the U.S. literature, and reasonably so: a tax that raises 0.1% to 0.3% of GDP can be replaced at almost any margin without visible consequences. At 0.5% of GDP and with financing feedbacks this large, the financing choice is first order.

One caveat applies to the comparison throughout. Labor supply is exogenous here, so raising τ_l has no efficiency cost beyond its income effect, while raising τ_k distorts the saving margin on which the reforms operate. The labor-tax results are therefore an upper bound on what that financing scheme would deliver with an elastic labor supply, and the gap between the two schemes is correspondingly an upper bound as well. The direction of the bias is clear even though its size is not: endogenous labor supply would narrow the gap rather than widen it.

6.4 Distributional Effects

Table 7 reports the effects of the reforms on the wealth distribution in the small open economy. The corresponding table for the case in which the capital income tax adjusts is in Appendix D.3 (Table A10).

	Gini	Top 1%	Top 5%	Top 10%	Top 20%
Benchmark ($\tau_b = 39.5\%$)	0.61	17.4	29.3	41.7	60.3
Eliminate ($\tau_b = 0$)	0.61	18.7 (+1.3)	30.0	42.1	60.5
Cut the rate ($\tau_b = 20\%$)	0.61	18.0 (+0.6)	29.5	41.8	60.3
Raise exemption (¥575M)	0.60	16.7 (−0.7)	28.5	41.2	60.1

Table 7: Distributional effects of inheritance tax reform in the small open economy, adjusting the labor income tax.

Notes: Gini coefficient of wealth and shares of total wealth held by households in the indicated top percentiles. Changes in the top 1% share relative to the benchmark, in percentage points, are in parentheses.

Three results stand out.

First, the share of wealth held by the richest households is decreasing in the inheritance tax rate. Eliminating the tax raises the top 1% share from 17.4% to 18.7%, an increase of 1.3 percentage points; cutting the rate to 20% raises it by 0.6 points. The effects extend down the distribution but attenuate quickly: the top 5% share rises by 0.7 points under elimination, the top 10% share by 0.4 points, and the top 20% share by 0.2 points.

Second, and in contrast, raising the exemption to ¥575 million *lowers* top-end concentration: the top 1% share falls by 0.7 percentage points and the Gini by 0.01. This is the most useful result of the section, because it means that base design and rate cuts are not substitutes distributionally even when they deliver similar aggregate effects.

The reason is that a higher exemption at an unchanged marginal rate relieves estates in near-inverse proportion to their size. Because the exemption shields a fixed amount, the *proportional* relief it delivers is largest for estates just above the old threshold and shrinks as the estate grows. An estate of ¥100 million, taxable in the benchmark, becomes untaxed outright. An estate of ¥5 billion sees its liability fall from $0.395 \times (5000 - 34.8)$ to $0.395 \times (5000 - 575)$ million yen, a reduction of about one tenth. A rate cut does the opposite: halving τ_b halves the liability on every taxable estate, whatever its size.

This is why the two reforms differ distributionally even when their aggregate effects are similar. The households relieved almost completely—those with estates between ¥34.8 and ¥575 million—sit in the upper-middle of the wealth distribution and are spread across a wide range of it, so their additional accumulation raises aggregate wealth without concentrating it. The households whose estates generate the top 1% share are relieved too, but only marginally, because for them the exemption is small relative to the estate. Aggregate wealth therefore rises faster below the very top than at it, and the top share falls.

It is worth being precise about one implication, since it is easy to state the point too strongly. The exemption reform is not a reform that spares the top. Appendix Table A15 shows that for the top 1% of the wealth distribution as a group the tax's share of wealth at death falls from 29.8% to about 6.2%, which is a larger reduction than halving the rate delivers. That is consistent with the above, because the top 1% of the wealth distribution is itself dominated by estates in the ¥34.8–575 million range; the estates for

which the exemption is immaterial are those of the super-high earnings state, whose mass is negligible but whose size drives the measured top 1% *share*. The reform relieves most of the rich substantially and the very richest hardly at all, and it is the second group that the concentration statistics track.

Third, the distributional effects are considerably more robust to the degree of capital mobility than the aggregate effects are. The same reforms in the closed economy (Table 9) move the top 1% share by 2.3 points under elimination against 1.3 points here, and by 0.8 points under the rate cut against 0.6—a difference of a factor of less than two, against a factor of close to nine for the aggregate wealth response to abolition. The exemption reform lowers the top 1% share in both cases. Whatever one concludes about how open Japanese capital markets are, the distributional ranking of the reforms is unaffected.

6.5 Importance of Parental Background Effects

In order to assess to what extent inheritance tax reform can affect the lifetime value of being born to a family with a given parental background, we recompute the one-time asset compensations of Section 5 for the economy in which the exemption is raised to ¥575 million while the rate is held at 39.5%. Table A7 in Appendix D.2 switches expected bequests only, keeping the child's initial productivity unchanged, while Table A8 switches both. In each table the panels report, in order, the net-bequest benchmark, the reform under net bequests, the gross-bequest benchmark, and the reform under gross bequests. The corresponding results when the capital income tax adjusts are in Appendix D.3.

Two results stand out, and the first runs counter to what the aggregate results would suggest.

The value of being born to super-high earning parents *falls* under the reform, even though the reform relieves large estates. The bequests-only compensation for moving from the 5th to the 1st background declines from 13.94 times average income in the benchmark to 13.63, and the total compensation from 18.03 to 17.68. Under the gross-bequest specification the decline is much smaller, from 13.91 to 13.85 and from 17.93 to 17.81.

The reason is the one that drives the distributional result of Section 6.4: these are compensations measured *relative* to other backgrounds, and a fixed exemption at an unchanged marginal rate relieves estates in near-inverse proportion to their size. The estates that generate the dynastic advantage of a 5th-state child are so large that a ¥575 million exemption is immaterial to them, while the estates of 3rd- and 4th-state parents are relieved outright. Those parents accumulate more and their children's prospects improve correspondingly, so the gap narrows. This is a consistency check on the distributional finding: the same reform that lowers the top 1% wealth share also lowers the measured value of a super-rich background, and for the same reason.

For everyone not born to a super-rich family, the value of parental background rises, but only slightly. The bequests-only compensation for moving from the 4th to the 1st background rises from 0.68 to 0.77 times average income, from the 3rd to the 1st from

0.37 to 0.38, and from the 2nd to the 1st from 0.04 to 0.07. These are movements of a few hundredths of average income. They reflect the fact that the newly exempted households are precisely the 3rd- and 4th-state parents whose estates sat between the old and the new exemption, so their children now expect a larger net inheritance while the children of 1st-state parents expect nothing either way.

A caveat, which we take up in Section 7.3: the numbers above are for the small open economy, and in the closed economy this particular result changes sign.

7 The Closed Economy, and the Comparison with the United States

The small open economy is our baseline case. We now report the quantitative results taking Japan as a closed economy, in which the net international investment position is held at zero, aggregate household assets equal the sum of domestic capital and government debt, and the return to capital adjusts to clear the capital market. This is the standard case in the literature, and in particular it is the one used by De Nardi and Yang (2016), which makes it the right basis for comparing our results with theirs.

7.1 Aggregate Effects

	τ_l	K	W	ΔK	ΔW	ΔY	ΔB	r_p (%)
Benchmark ($\tau_b = 39.5\%$)	0.248	2.733	3.940	—	—	—	—	3.76
Eliminate ($\tau_b = 0$)	0.253	2.794	4.010	+2.2%	+1.8%	+0.8%	−3.4%	3.64
Cut the rate ($\tau_b = 20\%$)	0.251	2.758	3.971	+0.9%	+0.8%	+0.3%	−1.9%	3.71
Raise exemption (¥575M)	0.252	2.761	3.975	+1.0%	+0.9%	+0.4%	−2.7%	3.71

Table 8: Aggregate effects of inheritance tax reform in the closed economy, adjusting the labor income tax.

Notes: K and W are the capital stock and aggregate household net worth, both relative to benchmark output. ΔK, ΔW, ΔY and ΔB are percentage changes in the aggregate capital stock, aggregate wealth, output and the flow of bequests relative to the benchmark. r_p is the portfolio return of (13). The net international investment position is zero in every row. Because $W = K + D$ and D/Y is fixed, aggregate wealth changes through both domestic capital accumulation and the change in government debt.

Three features of Table 8 deserve comment.

First, the household wealth response to abolition is approximately one-ninth of that in the small open economy. Domestic capital rises by 2.2% and output by 0.8%. Because government debt rises with output, household net worth increases by approximately 1.8%, compared with 15.3% in the small open economy. The response ratio is approximately 8.6. In the closed economy, additional saving is absorbed by domestic capital and government

debt rather than foreign assets. Capital deepening lowers the return to capital and reduces the portfolio return from 3.76% to 3.64%, substantially damping accumulation.

Second, the flow of bequests *falls*, by 3.4% under elimination, even though the capital stock rises. This may look paradoxical but follows from the general equilibrium price movements. Capital deepening raises the wage and lowers the interest rate, which shifts accumulation toward the young, who earn more, and away from the elderly, who earn less on their assets and decumulate faster. The additional capital is held disproportionately by households who are not about to die, so the annual flow released at death declines.

Third, the labor income tax must *rise*, from 24.8% to 25.3%, where in the small open economy it falls. With a much smaller increase in the wealth stock, the expansion of the capital income tax base no longer covers the inheritance tax revenue forgone.

7.2 Distributional Effects

Table 9 reports the effects on the wealth distribution. The comparison with Table 7 is instructive, because the distributional effects turn out to be much less sensitive to capital mobility than the aggregate effects, and because the one place they do differ is informative about the mechanism.

	Gini	Top 1%	Top 5%	Top 10%	Top 20%
Benchmark ($\tau_b = 39.5\%$)	0.61	17.4	29.3	41.7	60.3
Eliminate ($\tau_b = 0$)	0.63	19.7 (+2.3)	31.7	44.2	62.6
Cut the rate ($\tau_b = 20\%$)	0.61	18.2 (+0.8)	30.1	42.6	61.3
Raise exemption (¥575M)	0.61	17.1 (−0.3)	29.3	42.1	61.0

Table 9: Distributional effects of inheritance tax reform in the closed economy, adjusting the labor income tax.

Notes: Gini coefficient of wealth and shares of total wealth held by households in the indicated top percentiles. Changes in the top 1% share relative to the benchmark, in percentage points, are in parentheses. Benchmark bequest specification; the labor income tax balances the budget.

The ordering of the three reforms is the same as at fixed prices. Eliminating the tax raises the top 1% share, by 2.3 percentage points against 1.3 points in the small open economy, and cutting the rate raises it by 0.8 points against 0.6. The rate reforms thus move concentration somewhat more when prices adjust, because the lower return to capital compresses the accumulation of ordinary savers more than that of the very rich, whose saving is driven by the bequest motive. But the difference is a factor of under two, against a factor of nine for the aggregate wealth response. Whichever case one prefers, the distributional reading of the reforms is the same.

The exemption reform again lowers the top 1% share, by 0.3 percentage points against 0.7 at fixed prices, so the conclusion that base design and rate cuts work in opposite directions at the very top survives the change in the capital-market assumption. What does change is what happens just below the top. In the small open economy the exemption

reform lowers every share we report; in the closed economy it lowers the top 1% share, leaves the top 5% share unchanged at 29.3, and *raises* the top 10% and top 20% shares, by 0.4 and 0.7 percentage points, respectively.

This is the proportional-relief mechanism of Section 6.4 showing through more clearly. The households the exemption relieves outright are those with estates between ¥34.8 and ¥575 million, which sit in the upper-middle of the distribution; the households at the very top are barely relieved. At fixed prices aggregate wealth rises by enough (5.9%) that every reported share falls even as the upper-middle gains in absolute terms. In the closed economy aggregate wealth barely moves—the capital stock rises by 1.0%—so there is almost no scale effect to mask the composition effect, and the gain of the upper-middle shows up directly as a higher top 10% and top 20% share. The two cases therefore agree on the sign of the effect at the very top and disagree only on whether the upper-middle's gain is visible as a share.

7.3 The Value of Parental Background Under the Two Polar Cases

Section 6.5 reported that the exemption reform *lowers* the measured value of being born to super-high earning parents in the small open economy. That result does not survive the change in the capital-market assumption. The reversal is worth setting out here, both for its own sake and because the mechanism behind it is the one that will drive the reversal of the welfare verdict in Section 8 below.

Table 10 collects the compensation required to move a newborn from the super-high (5th) parental background to the lowest, across both cases, both financing instruments and both bequest specifications. The full matrices behind the closed-economy columns are in Appendix D.3.

		Small open economy		Closed economy	
	Benchmark	Labor tax	Capital tax	Labor tax	Capital tax
Bequests only					
Net bequests	13.94	13.63	13.58	14.15	14.18
Gross bequests	13.91	13.85	14.68	14.34	14.54
Bequests and human capital					
Net bequests	18.03	17.68	17.61	18.25	18.29
Gross bequests	17.93	17.81	18.70	18.41	18.62

Table 10: The value of a super-high parental background under the exemption reform, by capital-market assumption and financing instrument.

Notes: Each entry is the one-time asset compensation, as a multiple of average income, required to make a newborn born to a 5th-state parent indifferent to being born to a 1st-state parent. The "Benchmark" column is the economy before the reform, computed separately for each bequest specification. The remaining columns are the economy in which the exemption is raised to ¥575 million with the rate held at its benchmark value.

For the benchmark specification the sign depends entirely on whether capital markets are open. In the small open economy the compensation falls, from 13.94 to 13.63 under labor-tax financing on the bequests-only measure, and from 18.03 to 17.68 on the total. In the closed economy it rises, to 14.15 and 18.25. The movements are of similar size in either direction, about two percent of the benchmark value, and within each case the financing instrument barely matters.

The reason is the one that will reappear in the welfare results of Section 8. These are relative magnitudes, and what changes between the two cases is which group the reform helps more. At fixed prices the reform relieves the households with estates between ¥34.8 and ¥575 million, who sit in the upper-middle of the distribution, while the estates that generate the dynastic advantage of a 5th-state child still face the full 39.5% marginal rate. The middle gains more than the top, so the gap narrows. Once prices adjust, the extra accumulation depresses the interest rate, and a lower return to saving is a much heavier blow to a household whose position is almost entirely an asset position than to one living mostly on labor income. The top is cushioned by its larger inheritance; the middle is not. The gap widens.

Under the gross-bequest specification the compensation rises in three of the four cases, and is essentially flat in the fourth. That the reversal is confined to the benchmark specification is a reason for caution about this particular number, and we would not want to rest a policy conclusion on its sign. What we would rest a conclusion on is the magnitude: in every one of the eight cells the movement is under one unit of average income against a base of 14 to 19, so whichever way the sign falls, the exemption reform does very little to the value of a super-rich background. Within the model, bequests rather than inherited ability account for most of the advantage of the highest parental background, as Section 5 showed; but that background has negligible mass, and raising the

exemption to the current U.S. level does not materially change how much the advantage is worth.

7.4 Comparison with De Nardi and Yang (2016)

Table 11 places our closed-economy results beside those of De Nardi and Yang (2016). Two cautions apply to the comparison. The first is the one set out in Section 2: the U.S. tax in their model is not the U.S. tax of today, and all the U.S. figures in the table describe their calibration, not current policy. The second is that our two models share a common lineage and a deliberately similar structure, which is what makes the comparison informative, but they are calibrated to different data in different decades.

	De Nardi and Yang (2016) (United States)	This paper (Japan)
The tax as modeled		
Revenue / Y	0.33%	0.50%
Fraction of decedents taxed	1.9%	9.4%
Effective marginal rate	21.4%	39.5%
Effective exemption	\$756K	¥34.8M
The bequest flow		
B/W	0.88%	1.00%
Average tax rate T_b/B	12.1%	12.6%
Effects of abolition		
ΔK	+1.55%	+2.2%
Δ top 1% wealth share	+1.8pp	+2.3pp
Top 1% wealth share, benchmark	35.2	17.4

Table 11: Comparison with De Nardi and Yang (2016).

Notes: U.S. figures are from De Nardi and Yang (2016) and describe the U.S. estate tax as it stood in their calibration, not as it stands today; see Section 2. Their abolition figures are from the general equilibrium (closed-economy) experiment with the labor income tax adjusting, reported in their online appendix (Tables 6 and 7). Japanese figures are from the benchmark and the closed-economy abolition experiment, with the labor income tax adjusting, and are taken from Tables 8 and 9. T_b/B is inheritance tax revenue relative to the flow of bequests.

The comparison yields three observations. The first concerns the average tax rate on bequests. The two taxes claim almost exactly the same share of the bequest flow: an average rate of 12.1% in their calibration against 12.6% in ours. The Japanese tax is much larger by every measure of its size—revenue relative to output, the fraction of decedents it reaches, its effective marginal rate—and yet the fraction of what passes at death that it takes is essentially the same. The reason is concentration. The average rate equals the effective marginal rate times the share of the bequest flow that lies above the exemption. Wealth in their calibration is so concentrated—the top 1% hold 35.2% of it,

against 17.4% in ours—that about 57% of the flow lies above an exemption of roughly 13 times average income; in ours, only about 32% lies above an exemption of about 7 times average income. Japan's higher rate is offset by a smaller taxable share. What the Japanese tax's breadth buys is not a heavier aggregate burden on bequests but the same burden collected from five times as many decedents.

The second observation concerns accumulation. The aggregate effect of abolition in Japan is 2.2% against their 1.55%, a factor of about 1.4. Our separately recalibrated no-motive economy generates roughly two-thirds of the Japanese benchmark bequest flow relative to wealth, but that comparison does not identify the tax responsiveness of the benchmark's accidental component. Applied to their Table 1, the same comparison gives 66% for the U.S. (0.58% against 0.88%), so the accidental shares in the two calibrations are nearly identical, and the cross-country response difference cannot be attributed to a difference in intentional shares.

The third is distributional. The top 1% wealth share rises by 2.3 percentage points in our closed-economy abolition against 1.8 points in theirs. In absolute terms these are similar, but the benchmark shares differ enormously: 17.4% in our Japanese calibration against 35.2% in theirs. Measured proportionally, the Japanese distributional effect is therefore about two and a half times larger. This is where the breadth of the Japanese tax shows up. A tax that reaches a tenth of decedents is doing distributional work across much more of the upper half of the wealth distribution than one that reaches two percent, even when the two claim the same share of the aggregate flow.

Taken together, these comparisons show that the Japanese calibration produces a larger capital response to abolition and a larger proportional top-one-percent wealth-share response. The bequest-flow decomposition provides additional evidence on the capacity of life-cycle and precautionary saving to generate estates, but neither that decomposition nor the cross-country comparison establishes the intentional share as a sufficient statistic for reform responses. Those responses also depend on tax design, recipient behavior, financing, and equilibrium returns.

8 Welfare Effects

Our incomplete-markets framework generates, absent any policy or exogenous change, a stationary distribution of wealth. Following Conesa, Kitao, and Krueger (2009), we employ an ex-ante welfare criterion—before ability is realized—that measures the expected lifetime utility of a newborn worker in a stationary equilibrium, with the expectation taken over idiosyncratic shocks and parental background. We also report the expected lifetime utility of a newborn conditional on their initial productivity.

8.1 The Welfare Measure

For a household at state $x = (t, a, z, \tilde{y}, S_p)$, define the compensation $\Lambda(x)$ as the solution to

$$V^1\big(t,\, a - \Lambda(x)\,\bar{y},\, z,\, \tilde{y},\, S_p\big) = V^0\big(t,\, a,\, z,\, \tilde{y},\, S_p\big), \tag{16}$$

where $\bar{y}$ is average income (about ¥5 million; see Appendix B), the unit in which all compensations in the paper are expressed, and V is whichever value function applies at x—V_w for a worker whose parent is alive, V_w^I for a worker whose parent is deceased, or V_r for a retiree. Thus $\Lambda(x)$ is the one-time change in assets, expressed as a multiple of average income, that makes a household in the reformed economy exactly as well off as in the benchmark: a positive $\Lambda(x)$ is the amount it could give up and still be indifferent, a welfare gain from the reform, and a negative $\Lambda(x)$ is the amount it would need to receive, a welfare loss.

Aggregating $\Lambda(x)$ over the stationary distribution m^* of the benchmark economy gives

$$\Lambda^{\text{avg}} \;=\; \sum_x m^*(x)\,\Lambda(x), \qquad \text{Fraction gaining} \;=\; \sum_x m^*(x)\,\mathbf{1}\{\Lambda(x) > 0\}, \tag{17}$$

with the average gain of winners and the average loss of losers computed as the corresponding conditional means. Restricting m^* to $t = 1$ gives the newborn average. Every entry we report, including "Ex ante (all)", is a newborn average, and the "initial earnings" rows further condition on the newborn's initial productivity.

Two things should be said about this measure. It is an average of state-specific compensations, not the single transfer that equates expected lifetime utility across the two economies, and the two differ because the marginal value of assets varies across newborn states. The sign of the expected-utility criterion is nonetheless pinned down whenever the compensation has the same sign at every newborn state: since $\Lambda(x) > 0$ is equivalent to $V^1(x) > V^0(x)$, a uniformly positive Λ implies that the common transfer that equates expected lifetime utility is also positive. This applies to the small open economy below, where the fraction gaining is one. It does not apply in the closed economy, where about a fifth of newborns gain, and we report the average compensation there as what it is.

A second limitation applies to everything in this section. These are comparisons across stationary equilibria, and in the small open economy the reformed equilibrium is one in which households hold an additional 60% of annual output in foreign assets. A newborn in that equilibrium enjoys the income from those assets without having borne the consumption forgone in accumulating them. The small-open-economy gains are therefore an upper bound on what a transition would deliver, and the welfare of the generations alive when a reform is enacted is not something these numbers speak to. Computing the transition is the natural next step, and we return to it in Section 11.

8.2 Results

Table 12 reports the welfare effects of the exemption reform in both cases and under both financing instruments. We report welfare for this reform alone because it is the only one of the three whose aggregate and distributional effects point in the same direction, which is what makes the ex-ante comparison interpretable; for abolition and the rate cut, a larger aggregate wealth stock coincides with higher concentration at the top, and a single ex-ante number would combine the two without separating them. Reporting welfare for all three reforms would be a useful extension.

	Small open economy		Closed economy	
	Labor tax	Capital tax	Labor tax	Capital tax
Panel A: net bequests (benchmark)				
Ex ante (all)	+0.061	+0.093	−0.022	−0.041
Initial earnings, 1st	+0.009	+0.015	−0.007	−0.011
Initial earnings, 2nd	+0.028	+0.045	−0.016	−0.026
Initial earnings, 3rd	+0.073	+0.112	−0.028	−0.050
Initial earnings, 4th	+0.153	+0.228	−0.037	−0.084
Initial earnings, 5th	+15.76	+25.13	−5.65	−11.05
Fraction gaining	1.000	1.000	0.188	0.153
Winners' average gain	+0.061	+0.093	+0.042	+0.022
Losers' average loss	−0.001	—	−0.037	−0.053
Panel B: gross bequests				
Ex ante (all)	+0.118	−0.043	−0.024	−0.032
Initial earnings, 1st	+0.021	−0.008	−0.006	−0.006
Initial earnings, 2nd	+0.061	−0.023	−0.015	−0.017
Initial earnings, 3rd	+0.143	−0.051	−0.030	−0.038
Initial earnings, 4th	+0.274	−0.099	−0.049	−0.073
Initial earnings, 5th	+31.71	−12.26	−1.67	−7.93
Fraction gaining	1.000	0.109	0.171	0.142
Winners' average gain	+0.118	+0.018	+0.032	+0.023
Losers' average loss	—	−0.050	−0.036	−0.041

Table 12: Welfare effects of raising the exemption to the current U.S. level.

Notes: Entries are one-time asset compensations defined by (16), expressed as multiples of average income; a positive number is a welfare gain from moving from the benchmark to the reformed economy. "Ex ante (all)" is the measure computed under the veil of ignorance; the "initial earnings" rows condition on the newborn's initial earnings state. "Fraction gaining" is the share of newborns with $\Lambda > 0$, and the last two rows are the average gain among those who gain and the average loss among those who lose. A dash in those rows means the set is of negligible measure. The reform holds τ_b at its benchmark value and raises x_b to ¥575 million.

In the small open economy, every newborn gains under the benchmark specification.

The ex-ante gain is 0.061 times average income when the labor income tax adjusts and 0.093 when the capital income tax adjusts, and the gains rise steeply and smoothly with the initial earnings draw—0.009, 0.028, 0.073 and 0.153 times average income in the first through fourth states, and 15.76 times average income for the vanishingly rare newborns who start in the super-high state, who are typically the children of super-high earners expecting the largest inheritances and whose expected inheritances rise by the most in absolute terms, even though the exemption relieves their very large estates only slightly in proportional terms. With prices fixed, the reform is in effect a transfer toward households connected to taxable estates, and its cost is not large enough to make any group worse off on average: the fraction gaining is one, and the average loss among the measure-zero set of losers is a thousandth of average income. That the gain is larger under capital-tax financing follows directly from Table 6: that is the scheme under which τ_k falls.

In the closed economy the verdict reverses, and it reverses for every earnings state at once. The ex-ante effect is a loss of 0.022 times average income under labor-tax financing and 0.041 under capital-tax financing, and fewer than one newborn in five gains—18.8% and 15.3% respectively. Every conditional entry is negative, including the super-high state, which loses 5.65 times average income under labor-tax financing and 11.05 under capital-tax financing where at fixed prices it had gained 15.76 and 25.13. The reason is the capital deepening the reform induces. In the closed economy the additional saving must be absorbed domestically, so the interest rate falls, and the loss from a lower return to saving falls on precisely the households whose accumulation the reform was meant to encourage—most heavily on the dynastically wealthy, whose entire position is an asset position. The wage rises, which helps, but not by enough for most newborns.

The last two rows of each panel make the point in a different way. In the closed economy roughly a fifth of newborns still gain, and those who do gain an average of 0.042 times average income under labor-tax financing, while the four-fifths who lose do so by an average of 0.037. The reform is not a wash for everyone; it is a small gain for a minority and a comparable small loss for the majority, with the ex-ante average slightly negative. The contrast with the small open economy, where the fraction gaining is one, is as sharp as the contrast in the aggregates.

Panel B puts the bequest specification alongside the capital-market assumption, and the pattern is instructive. In the closed economy the specification does not matter for the verdict: every closed-economy entry in Panel B is negative, under both bequest motives and both financing instruments, exactly as in Panel A. It matters only for the magnitude: the super-high state loses 1.67 times average income under gross bequests with labor-tax financing against 5.65 under net bequests.

At fixed prices, by contrast, the specification can change the sign. Under gross bequests with labor-tax financing the gains are roughly twice those of the benchmark, which is what the larger aggregate response reported in Section 9 below would suggest; but under gross bequests with capital-tax financing the ex-ante effect turns negative, -0.043, and only 11% of newborns gain. That is the single cell in the whole table whose sign turns on the bequest motive rather than on capital mobility, and a likely reason is that the gross-

bequest calibration requires a higher effective rate, 45.9%, so the reform forgoes more revenue and the required capital income tax increase is correspondingly larger. We would not build an argument on it. The robust reading of Table 12 is that capital mobility governs the sign and the bequest specification governs the magnitude.

The sign of the welfare verdict, then, depends on how open Japanese capital markets are. This is a sharper version of the point made in Section 6.3: in an economy where the inheritance tax is large enough to move aggregates, conclusions about its reform are hostage to assumptions that are of second-order importance where the tax is small.

It is worth noting how this relates to De Nardi and Yang (2016). They find that *raising* U.S. estate taxation to the statutory levels then in force generates ex-ante welfare gains for a newborn, financed by very large losses for the super-rich, with the gains operating mainly through the redistribution of a small amount of revenue. Our closed-economy result moves in the same direction: narrowing the base lowers ex-ante welfare. The two findings are consistent, and the reason they can be stated so much more sharply for Japan is that the general equilibrium price effects, which are negligible in their setting, are first order in ours.

9 Alternative Bequest Motives

All of the results above use the benchmark specification, in which parents value the estate net of inheritance tax. This section turns to the two alternatives: the gross-bequest specification, in which parents value the pre-tax estate, and the specification with no voluntary motive at all. We first use the latter to decompose the bequest flow, and then compare reform responses across the three specifications.

9.1 A Model-Based Decomposition of the Bequest Flow

Section 5 compares the benchmark with an economy in which the voluntary bequest motive is absent. That comparison also answers a question that matters for interpreting the reforms: how much of what passes at death in Japan is deliberately accumulated in order to be handed on, and how much is simply wealth that households happened to be holding when they died.

We answer it with the decomposition of Hendricks (2002). The economy without a warm glow is recalibrated to reproduce the same two administrative facts about the inheritance tax, so the bequests it still generates are purely the by-product of mortality acting on life-cycle and precautionary saving. We label that component *accidental*; the difference between the benchmark flow and it is *intentional*. Table 13 reports the result.

	B/W
Benchmark (net bequests)	1.00%
No voluntary bequest motive	0.67%
Implied accidental share	$\approx 67\%$
Implied intentional share	$\approx 33\%$

Table 13: The accidental share of the Japanese bequest flow.

Notes: B/W is the annual flow of bequests relative to aggregate household net worth. Both economies are calibrated to the same inheritance tax revenue relative to output and the same fraction of decedents leaving taxable estates; see Table 2. Following Hendricks (2002), the flow generated without a voluntary motive is interpreted as the accidental component.

Under this decomposition convention, the recalibrated economy without a voluntary motive generates approximately two-thirds of the benchmark bequest flow relative to wealth; the remaining third is the residual associated with introducing the motive and recalibrating. This is a comparison across specifications, not a household-level partition of benchmark estates into responsive and unresponsive components. Without a warm glow, the inheritance tax no longer directly changes the donor's utility from leaving an estate. It nevertheless changes heirs' net receipts and saving, with further effects through portfolio returns and, in the closed economy, factor prices. The decomposition therefore does not establish that only the intentional third responds to reform.

For the United States, Hendricks (2002) finds that accidental bequests account for at least half of the flow and possibly all of it. Our estimate for Japan is in the same range, and the same comparison in De Nardi and Yang (2016), whose model is closest to ours, gives 66% for the U.S.: their recalibrated economy without a voluntary motive generates a bequest–wealth ratio of 0.58%, against 0.88% in their benchmark.

We also examine the equilibrium responsiveness of the inheritance tax base. Let $\mathcal{B}_b = T_b/\tau_b$ denote taxable estate value above the exemption, where T_b is inheritance tax revenue and τ_b is the model's effective tax rate. For the cut from each specification's benchmark rate to 20%, we compute the log arc elasticity

$$\eta_b = \frac{\log(\mathcal{B}_b^1/\mathcal{B}_b^0)}{\log[(1-\tau_b^1)/(1-\tau_b^0)]}.$$

Across both bequest specifications, both financing instruments and both cases, the elasticities range from -0.15 to 0.56. The largest positive responses occur in the small open economy with labor-tax financing: 0.43 under net bequests and 0.56 under gross bequests. In the other cases, they range from -0.15 to 0.08. All are inelastic in the conventional sense, but their variation shows that the equilibrium response depends on financing and returns. These calculations do not isolate a direct donor response or identify the elasticity of the benchmark's accidental component.

The return movements help interpret this pattern. In the small open economy with labor-tax financing, the rate cut raises wealth, reduces the share of zero-interest debt in

household portfolios, and increases the after-tax portfolio return. This reinforces accumulation and is associated with the largest positive tax-base elasticities. Under capital-tax financing, the higher capital income tax lowers the after-tax return; in the closed economy, capital deepening lowers the return to capital. These offsetting forces are associated with small or negative base responses. The comparisons are consistent with an important role for return feedback, but do not separately quantify the direct donor, recipient, financing, and return channels.

9.2 Reform Responses Under Alternative Bequest Motives

Table 14 reports the two reforms of greatest interest under the two alternatives. Each is separately recalibrated to the moments of Table 2 that apply to it, as described in Section 4.

	τ_b	x_b	Eliminate	Exemption ¥575M
Small open economy: ΔW				
Net bequests (benchmark)	39.5%	¥34.8M	+15.3%	+5.9%
Gross bequests	45.9%	¥36.2M	+17.4%	+9.6%
No voluntary bequests	62.1%	¥25.0M	+6.5%	+4.3%
Closed economy: ΔK				
Net bequests (benchmark)	39.5%	¥34.8M	+2.2%	+1.0%
Gross bequests	45.9%	¥36.2M	+2.8%	+1.3%
No voluntary bequests	62.1%	¥25.0M	+2.9%	+2.0%

Table 14: Inheritance tax reform under alternative bequest motives.

Notes: ΔW is the percentage change in aggregate household net worth; ΔK the percentage change in the capital stock. The labor income tax adjusts to balance the budget in all cases. Each specification is separately recalibrated to the moments of Table 2 that apply to it, which is why τ_b and x_b differ across rows.

Three points emerge.

First, comparing net with gross bequests, the aggregate responses are similar, if somewhat larger under gross bequests. One offsetting mechanism is that parents who value the after-tax estate can increase gross saving when taxation rises in order to protect their heirs' receipts. Lower taxation withdraws that incentive alongside reducing the marginal tax wedge. This mechanism is absent when parents value the gross estate. The tax-base elasticities for the rate cut are also higher under gross bequests, as reported in Section 9.1. These patterns are consistent with an offsetting donor response under net-bequest preferences, but the separately calibrated specifications and equilibrium feedbacks prevent an isolated quantitative attribution.

Second, the recalibrated economy without a voluntary motive has a smaller small-open-economy response to abolition: 6.5% against 15.3% in the benchmark. Its response reflects changes in heirs' net receipts, fiscal financing, and endogenous portfolio returns

despite the absence of a direct bequest-utility channel. Because the specifications differ in patience and effective tax parameters, their response gap does not isolate the voluntary motive's contribution or independently confirm the 67% flow decomposition.

Third, in the closed economy the three specifications are much closer together, and the ordering is even reversed: the no-motive specification produces a slightly larger capital response, which underscores the distinction between flow composition and reform responsiveness. This is what one would expect: there the general equilibrium price response does most of the work, and that response is common across specifications, so the differences in the strength of the bequest motive are largely washed out.

10 Population Aging

Japan's population will continue to age for decades, and the common presumption in the policy debate is that this strengthens the case for taxing inheritances: the elderly hold a large and growing stock of wealth, a "great wealth transfer" is coming, and a tax on that transfer is an attractive and growing revenue source. This section asks what the model says.

We replace the survival probabilities of the 1990 birth cohort with those projected for the 2070 birth cohort by the Institute of Population and Social Security Research (Figure A2 in the appendix compares the two schedules), hold preferences and the inheritance tax parameters (τ_b, x_b) fixed, and let the payroll tax τ_s and the labor income tax τ_l rebalance the government budget. The discount factor is not recalibrated, so the exercise answers what happens to an economy with these preferences when it faces 2070 mortality, rather than what a model calibrated to 2070 Japan would look like.

10.1 The Benchmark Economy Under 2070 Mortality

	L	K	W	W/Y	B	B/W	B/Y	T_b/Y
Baseline mortality	0.695	2.733	3.940	3.940	0.0395	1.003%	3.953%	0.499%
2070 projections	0.671	2.640	3.889	4.025	0.0364	0.935%	3.763%	0.449%
Change	−3.4%	−3.4%	−1.3%	+2.1%	−8.0%	−6.8%	−4.8%	−10.0%

Table 15: The benchmark economy under baseline and 2070 mortality.

Notes: L is aggregate effective labor, W aggregate household net worth, B the annual flow of bequests and T_b inheritance tax revenue. K, W and B are relative to baseline output, which is therefore one in the first row. With the capital–labor ratio pinned, K and Y fall in the same proportion, so $K/Y = 2.733$ in both rows, and output falls to 0.966 under the 2070 projections, which is why B/Y falls by less than B. The payroll tax and the labor income tax adjust to balance the government budget. Small open economy, benchmark bequest specification.

The mechanics are worth setting out in order, because three of the entries in Table 15 move in directions that are easy to get wrong.

Aggregate effective labor falls by 3.4%. More of each cohort survives past the mandatory retirement age of 67, so a larger share of the population is retired and a smaller share is working. With the capital–labor ratio pinned by the world interest rate, the capital stock and output fall in the same proportion.

Wealth relative to output rises by 2.1%, for two reasons that both push in the same direction. A longer retirement must be self-financed, so households hold more wealth at every age; and the denominator shrinks with the labor force. Note, though, that the *level* of wealth falls slightly, by 1.3%: the composition effect on W/Y comes mostly from the denominator.

The flow of bequests nonetheless falls, by 8.0% in level and 6.8% relative to wealth. This is the central result of the section. Part of the fall comes from fewer estates. With zero population growth, annual deaths equal annual births, but longer lives raise the population each cohort supports, so deaths per person fall by 3.4%, from 1.60% to 1.55% of the population each year. The rest comes from smaller estates: deaths shift to older ages, at which households have drawn their wealth down further, so the average estate falls by about 4.8%. Aggregate bequests fall by 8.0% while aggregate wealth falls by only 1.3%. Aging thus enlarges the wealth stock relative to output while shrinking the flow released at death, and for a tax levied on the flow rather than the stock, the second effect is what counts. Inheritance tax revenue per unit of output falls by 10%, even though wealth relative to output is 2% higher. Meanwhile τ_l must rise from 24.8% to 25.3% to help fund the larger pension burden.

10.2 Reform Responses Under 2070 Mortality

	ΔW (small open)		ΔK (closed)		Δ top 1% (pp)	
	Baseline	2070	Baseline	2070	Baseline	2070
Net bequests (benchmark)	+15.3%	+7.7%	+2.2%	+1.7%	+1.3	+0.9
Gross bequests	+17.4%	+13.1%	+2.8%	+2.7%	+2.6	+2.9
No voluntary bequests	+6.5%	+4.1%	+2.9%	+1.9%	+1.6	+1.1

Table 16: Effects of eliminating the inheritance tax under baseline and 2070 mortality.

Notes: All entries are for the reform that eliminates the inheritance tax, with the labor income tax adjusting. ΔW and ΔK are percentage changes in household net worth and in the capital stock; Δ top 1% is the change in the top 1% wealth share in percentage points, in the small open economy.

Almost every response shrinks. In the benchmark specification the small open economy response to eliminating the tax falls from 15.3% to 7.7%, roughly by half; the closed-economy response falls from 2.2% to 1.7%; and the increase in the top 1% share falls from 1.3 to 0.9 percentage points. The one exception in Table 16 is the distributional effect

under gross bequests: the increase in the top 1% wealth share rises slightly, from 2.6 to 2.9 percentage points; every other entry in the table falls.

Longer lives change the composition of bequests. Applying the decomposition of Section 9.1 under 2070 mortality, the bequest–wealth ratio without a voluntary motive falls from 0.67% to 0.55%, against 1.00% to 0.94% in the benchmark with net bequest motive, so the accidental share falls from 67% to 58%. Relative to wealth, accidental bequests fall by 19% and intentional bequests rise by 18%.

Both movements follow from the change in survival. A household saving purely for life-cycle and precautionary reasons, given the two to two and a half additional years of expected retirement that the 2070 schedule implies, is much better at spending its wealth down: longer life is precisely what allows people to *avoid* dying with money left over. A household with a warm-glow motive, in contrast, must fund a longer retirement *and* still wants to leave something behind, so it holds more wealth at every age. The decline in accidental bequests dominates: relative to wealth, they fall by 0.13 percentage points while intentional bequests rise by 0.06 points, so the total falls even as the intentional part grows.

Because intentional bequests respond directly to the tax through the donor's bequest utility, while accidental bequests do not, a larger intentional share would on its own make the reform responses larger. They nevertheless shrink, for two reasons. The first is that there is less to give back. Inheritance tax revenue per unit of output is 10% smaller under 2070 mortality, so eliminating the tax is a smaller transfer and has a smaller income effect. The second is that the bequest arrives later. A warm-glow payoff further in the future is discounted more heavily at the margin, so a given cut in τ_b buys less additional saving today; the same is true on the receiving side, since the heir inherits later in life, after having already accumulated more on their own account and with fewer years left in which to use the inheritance.

10.3 What Aging Means for the Policy Debate

The common presumption with which this section opened confuses the stock with the flow. An inheritance tax reaches wealth only when it changes hands. Aging raises the stock relative to output but slows its turnover, and for a transfer tax the second effect dominates: revenue per unit of output falls by 10% even as W/Y rises by 2%.

The revenue case for the tax therefore weakens as the population ages, because its base shrinks relative to output even as the wealth stock grows. The welfare case for cutting it weakens as well. Table 17 repeats the welfare calculation of Section 8 for the exemption reform under 2070 mortality. In the small open economy with labor-tax financing, the ex-ante gain falls from 0.061 to 0.003 times average income, and the share of newborns who gain falls from all of them to 36%; with capital-tax financing, the gain of 0.093 becomes a loss of 0.073. In the closed economy, where the reform already lowers ex-ante welfare, the loss deepens slightly. Aging thus lowers the stakes on both sides of the debate.

	Small open economy		Closed economy	
	Labor tax	Capital tax	Labor tax	Capital tax
Ex ante (all)				
Baseline mortality	+0.061	+0.093	−0.022	−0.041
2070 projections	+0.003	−0.073	−0.036	−0.043
Fraction gaining				
Baseline mortality	1.000	1.000	0.188	0.153
2070 projections	0.363	0.011	0.152	0.110

Table 17: Welfare effects of the exemption reform under baseline and 2070 mortality.

Notes: Net-bequest specification. Entries are defined as in Table 12: "Ex ante (all)" is the one-time asset compensation, in multiples of average income, computed under the veil of ignorance, and "Fraction gaining" is the share of newborns who gain. Under gross bequests the small open economy with labor-tax financing still gains (0.070 under 2070 mortality, against 0.118).

One caveat is important. This is a comparison of two steady states, and the transition between them runs the other way in the near term. As today's unusually large elderly cohorts die, the annual flow of bequests will first *rise*, perhaps substantially, before settling at the lower level that the 2070 mortality rates imply. The near-term and long-run revenue outlooks therefore point in opposite directions, and nothing in the steady-state comparison here should be read as a forecast of inheritance tax revenue over the next two decades. Computing that transition is the most important extension of this paper.

11 Conclusions and Directions for Future Research

Japan taxes inheritances more heavily than almost any other advanced economy, and taxes them far more broadly: nearly a tenth of Japanese decedents leave a taxable estate, against about one in a thousand in the United States. This paper provides, to our knowledge, the first quantitative evaluation of reforming it in general equilibrium. We build an overlapping generations, incomplete-markets model with idiosyncratic earnings risk, voluntary and accidental bequests, and intergenerational transmission of earnings ability, adapted to Japanese demographics, portfolios and fiscal institutions, and calibrated to Japanese earnings and wealth inequality together with the observed revenue, breadth and concentration of taxable estates. The calibrated effective exemption of ¥34.8 million and effective rate of 39.5% line up closely with the statutory tax code, although no feature of that code is targeted, and we regard this as an important validation of the framework.

Five findings stand out.

First, two modeling choices each move the answer by an order of magnitude: where the marginal unit of saving goes, and which tax replaces the revenue. Eliminating the inheritance tax raises household net worth by 15.3% in the small open economy and by approximately 1.8% in the closed economy, a roughly ninefold response ratio; in the closed

economy, domestic capital rises by 2.2% and output by 0.8%. Switching the balancing instrument from the labor income tax to the capital income tax moves the same response from 15.3% to 1.2%, a comparable factor. Both dominate the choice between net and gross bequest utility, which moves the response by less than a third. Japan's small share of world output makes the small open economy the closer approximation; the closed economy shows how much changes if additional saving must be absorbed at home. Within the open economy, about two-thirds of the response operates through the endogenous portfolio return: as household net worth grows, the zero-yielding government debt that makes up 31% of Japanese portfolios is diluted, the average return to saving rises, and that reinforces the direct effect of the reform. This channel has no counterpart in models without government debt in household portfolios.

Second, substantial bequest flows arise without a voluntary motive. Under the model-based decomposition of Hendricks (2002), the separately recalibrated no-motive economy generates roughly two-thirds of the benchmark bequest flow relative to wealth. This comparison does not identify a tax-unresponsive component within the benchmark: heirs' receipts, saving, financing taxes, and equilibrium returns can respond even without a warm glow. In the cross-country comparison, the Japanese and modeled U.S. taxes claim similar shares of the bequest flow, 12.6% and 12.1%, while Japan's capital response to abolition is about 1.4 times larger. These comparisons document equilibrium outcomes rather than isolate bequest composition as their cause.

Third, the design of the base and the level of the rate are not substitutes. Raising the exemption to the current U.S. level delivers about two-fifths of the aggregate gain of abolition while *lowering* the top 1% wealth share by 0.7 percentage points; halving the effective rate to 20% delivers a similar aggregate gain but raises the top share. The reason is that a fixed exemption at an unchanged marginal rate relieves estates in near-inverse proportion to their size—almost completely just above the old threshold, hardly at all for an estate of several billion yen—whereas a rate cut relieves every taxable estate in the same proportion. Since it is the largest estates that the concentration measures track, the two instruments move the top share in opposite directions even when their aggregate effects are similar.

Fourth, two choices that are of second-order importance where the tax is small can each reverse a conclusion where it is large. Under capital-income-tax financing, the exemption reform becomes the largest of the three reforms rather than the smallest, because it is the only one that expands the capital income tax base by more than the revenue it forgoes. And the welfare verdict flips between the two cases: at fixed prices every newborn gains from the exemption reform, while in the closed economy it lowers ex-ante welfare and fewer than one newborn in five gains, because the capital deepening it induces lowers the return to saving and that loss falls on precisely the households the reform was meant to encourage.[14] The measured value of being born to super-high earning parents moves the same way, falling slightly at fixed prices and rising slightly once prices adjust, though in

[14]The welfare gains in the small open economy depend on the assumed return and tax treatment of the increase in NIIP.

that case the movement is small enough in either direction—under one unit of average income against a base of 14 to 19—that the reform should be read as leaving dynastic advantage essentially where it was.

Fifth, population aging shrinks almost every response, roughly by half in the baseline. Longer lives lower the accidental share of bequests from 67% to 58% under our model-based decomposition, which on its own would make bequests more responsive to the tax; responses shrink nonetheless because the flow of bequests shrinks relative to the wealth stock and arrives later in the recipient's life. The policy implication is that the common presumption in the debate, that an aging Japan makes the inheritance tax a more attractive revenue source, confuses the stock with the flow: revenue per unit of output falls by 10% even as wealth relative to output rises by 2%. Aging also shrinks the welfare gain from cutting the tax: in the small open economy the newborn gain from the exemption reform nearly vanishes. Aging thus weakens both the revenue case for the tax and the welfare case for cutting it, lowering the stakes on both sides of the debate.

Our framework makes several assumptions that suggest directions for future research. First, and most important, we limit ourselves to steady-state comparisons. As Section 10 notes, the transition runs the other way in the near term: as today's unusually large elderly cohorts die, the flow of bequests will first rise before settling at its lower long-run level, so the near-term and long-run revenue outlooks point in opposite directions. Second, we assume exogenous labor supply, and thus abstract from the labor supply distortions of the labor income tax adjustments that finance the reforms; endogenizing labor supply would tend to make base-narrowing reforms financed by labor taxes less attractive. Third, we abstract from the composition of estates and from avoidance: in Japan, real estate receives preferential valuation and there are special provisions for business succession, so modeling the portfolio choice between housing and financial assets, and the costs of avoidance, would sharpen the mapping between statutory and effective taxation. Fourth, we do not model gift taxation and inter vivos transfers, which in Japan are integrated with the inheritance tax and are themselves the object of ongoing reform. Finally, extending the analysis to incorporate entrepreneurship, in the spirit of Cagetti and De Nardi (2009) and Guo (2026), would allow us to address the business-succession concerns that figure prominently in the Japanese policy debate.

References

Attanasio, O. P. and G. Weber (1993). Consumption growth, the interest rate and aggregation. *The Review of Economic Studies 60*(3), 631–649.

Bosworth, B. P. and S. Anders (2011). Saving and wealth accumulation in the PSID, 1984–2005. NBER Working Paper 17689.

Cagetti, M. and M. De Nardi (2009). Estate taxation, entrepreneurship, and wealth. *The American Economic Review 99*(1), 85–111.

Cashin, D. and T. Unayama (2016). Measuring intertemporal substitution in consumption: Evidence from a VAT increase in Japan. *The Review of Economics and Statistics 98*(2), 285–297.

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003). Accounting for U.S. earnings and wealth inequality. *Journal of Political Economy 111*(4), 818–857.

Conesa, J. C., S. Kitao, and D. Krueger (2009). Taxing capital? Not a bad idea after all! *American Economic Review 99*, 25–48.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies 71*(3), 743–768.

De Nardi, M., S. İmrohoroğlu, and T. J. Sargent (1999). Projected U.S. demographics and social security. *Review of Economic Dynamics 2*(3), 575–615.

De Nardi, M. and F. Yang (2016). Wealth inequality, family background, and estate taxation. *Journal of Monetary Economics 77*, 130–145. Carnegie-Rochester-NYU Conference Series on Public Policy.

Fuster, L., A. İmrohoroğlu, and S. İmrohoroğlu (2003). A welfare analysis of social security in a dynastic framework. *International Economic Review 44*(4), 1247–1274.

Fuster, L., A. İmrohoroğlu, and S. İmrohoroğlu (2007). Elimination of social security in a dynastic framework. *Review of Economic Studies 74*(1), 113–145.

Fuster, L., A. İmrohoroğlu, and S. İmrohoroğlu (2008). Altruism, incomplete markets, and tax reform. *Journal of Monetary Economics 55*(1), 65–90.

Guo, Y. (2026). Inheritance, entrepreneurship, and estate taxation. *Journal of Economic Dynamics and Control 189*, 105359.

Hansen, G. and S. İmrohoroğlu (2023). Demographic change, government debt and fiscal sustainability in Japan: The impact of bond purchase by the Bank of Japan. *Review of Economic Dynamics 50*, 88–105.

Hayashi, F. and E. C. Prescott (2002, January). The 1990s in Japan: A lost decade. *Review of Economic Dynamics 5*(1), 206–235.

Hendricks, L. (2002). Intended and accidental bequests in a life-cycle economy. Working paper, SSRN 304721, available at `https://ssrn.com/abstract=304721`.

Horioka, C. Y. (2014). Are Americans and Indians more altruistic than the Japanese and Chinese? Evidence from a new international survey of bequest plans. *Review of Economics of the Household 12*(3), 411–437.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (1995). A life cycle analysis of social security. *Economic Theory 6*(1), 83–114.

Kitao, S., M. Suzuki, and T. Yamada (2025). Earnings inequality analysis using administrative tax data. RIETI Discussion Paper Series 25-J-020, Research Institute of Economy, Trade and Industry (RIETI). (in Japanese).

Kitao, S. and T. Yamada (2025). Earnings, income, and wealth inequality in Japan: a long-term perspective, 1984–2019. *The Japanese Economic Review 76*, 231–283.

National Tax Agency (2025). The 149th National Tax Agency Annual Statistics Report, FY2023. Technical report, National Tax Agency of Japan, Tokyo.

Niimi, Y. and C. Y. Horioka (2019). The wealth decumulation behavior of the retired elderly in Japan: The relative importance of precautionary saving and bequest motives. *Journal of the Japanese and International Economies 51*, 52–63.

OECD (2021). *Inheritance Taxation in OECD Countries*. Number 28 in OECD Tax Policy Studies. Paris: OECD Publishing.

OECD (2025). *Pensions at a Glance 2025: OECD and G20 Indicators*. Paris: OECD Publishing.

Ono, T. (2024). Bequests and wealth inequality in Japan. *Journal of the Japanese and International Economies 72*, 101314.

Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.

Ueda, A. (2015, September). An empirical analysis on intergenerational persistence of income in Japan. WINPEC Working Paper Series E1511, Waseda Institute of Political Economy (WINPEC), Waseda University.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy 110*(4), 825–853.

Yamada, T. (2025, December). Macroeconomic and fiscal implications of changes in the labor share under population aging in Japan. *Journal of the Economics of Ageing 32*, 100608.

Yang, F. (2013). Social security reform with impure intergenerational altruism. *Journal of Economic Dynamics and Control 37*(1), 52–67.

Yogo, M. (2004). Estimating the elasticity of intertemporal substitution when instruments are weak. *The Review of Economics and Statistics 86*(3), 797–810.

A Model Demographics and Stationary Equilibrium

A.1 Model Demographics

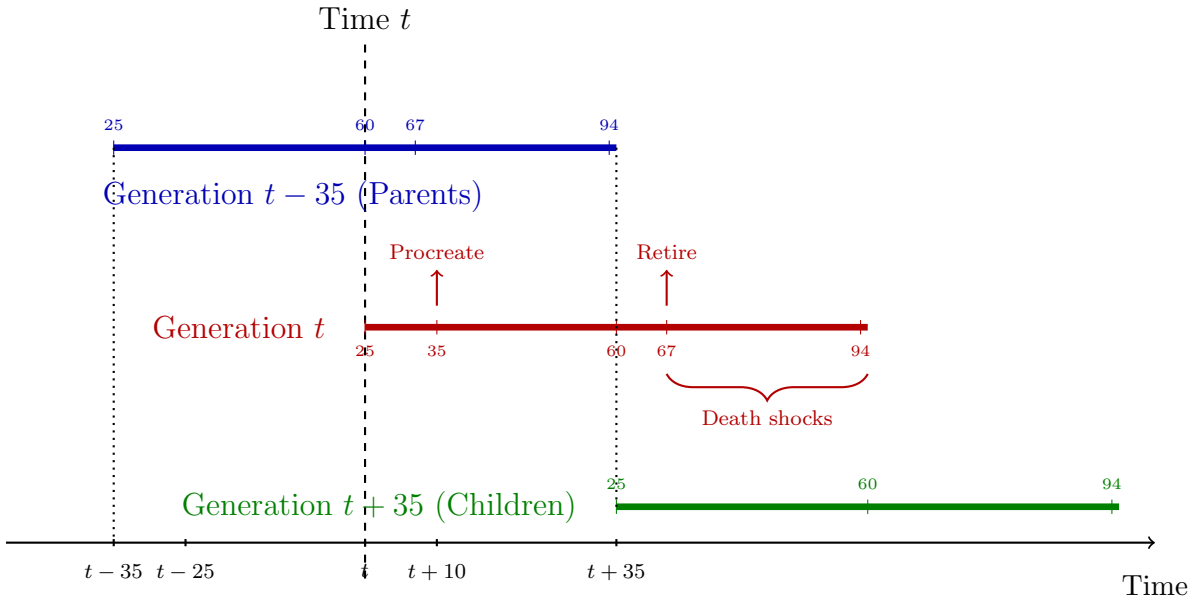


Figure A1: Model Demographics.

A.2 Definition of Stationary Equilibrium

We focus on a stationary equilibrium in which factor prices and the joint distribution of age, wealth, and other household states are constant over time. Let X denote the household state space and $x \in X$ a household state. Population growth is zero.

A stationary equilibrium consists of government tax rates, pension benefits, spending, and debt $(\tau_l, \tau_s, \tau_k, \tau_d, \tau_b, x_b, P(\widetilde{y}), G, D)$; a capital return r_k, a wage rate w, and a common household debt-holding requirement θ; value functions $V(x)$ and household allocations $c(x)$ and $a'(x)$; aggregate consumption C, household wealth W, domestic capital K, and effective labor L; and an invariant household distribution m^*. In the small open economy, the equilibrium also includes a net foreign asset position NIIP and a trade balance TB. These objects satisfy the following conditions.

(i) Household optimality. Given tax rates, pension benefits, asset returns, the wage, and the debt-holding requirement θ, the functions $V(x)$, $c(x)$, and $a'(x)$ solve the household problems specified in Section 3.5.

(ii) Distributional stationarity. The distribution m^* is invariant under the household decision rules and the exogenous demographic and earnings transitions.[15]

[15] We normalize $m^*(X) = 1$. For any measurable subset $B \subseteq X$, $m^*(B)$ is the fraction of the living

(iii) Aggregation and market clearing. Aggregate consumption, wealth, and effective labor satisfy

$$C = \int_X c(x)\, m^*(dx),$$
$$W = \int_X a\, m^*(dx),$$
$$L = \int_X I_{t\le 42}\, y\, m^*(dx),$$

where y is total worker productivity, including the age-efficiency profile.

Model production and government debt satisfy

$$Y = F(K, L), \qquad D = \bar{d}Y,$$

where $\bar{d}$ is held fixed across equilibria. The common household debt-holding requirement is

$$\theta = \frac{D}{W}.$$

In the closed economy, asset-market clearing and the aggregate resource constraint are

$$W = K + D,$$
$$C + \delta K + G = F(K, L),$$

with NIIP $= 0$.

In the small open economy, asset-market clearing becomes

$$W = K + D + \text{NIIP}.$$

Foreign assets earn the same pre-tax net return as domestic capital, $\bar{r}_k$. The stationary resource and external-account conditions are

$$C + \delta K + G = F(K, L) + \bar{r}_k\, \text{NIIP},$$
$$\text{TB} \equiv F(K, L) - C - \delta K - G,$$
$$\text{TB} + \bar{r}_k\, \text{NIIP} = 0.$$

Thus, in a stationary equilibrium, income from a positive net foreign asset position finances an equal-sized trade deficit, leaving the current account balanced and the net foreign asset position unchanged. The reference benchmark has zero NIIP; reforms can change its stationary level.

(iv) Factor pricing. In the closed economy, factor prices satisfy

$$r_k = F_K(K, L) - \delta, \qquad w = F_L(K, L).$$

In the small open economy, the capital return is fixed at the return in the calibrated

population whose states belong to B.

reference equilibrium, $\bar{r}_k$. Domestic capital and the wage satisfy

$$r_k = \bar{r}_k, \qquad F_K(K, L) - \delta = \bar{r}_k, \qquad w = F_L(K, L).$$

(v) Government budget balance. The general government budget satisfies

$$\begin{aligned} G + r_d D = \; & [(1-\theta)\tau_k r_k + \theta \tau_d r_d]\, W + \tau_l w L \\ & + \tau_b \int_X (1 - p_t) I_{t>42} \max\{a'(x) - x_b, 0\}\, m^*(dx) \\ & + \tau_l \int_X I_{t>42} P(\widetilde{y})\, m^*(dx). \end{aligned}$$

The public pension budget is separately balanced:

$$\int_X I_{t>42} P(\widetilde{y})\, m^*(dx) = \tau_s \int_X I_{t\leq 42} \min\{wy, \widetilde{y}_c\}\, m^*(dx).$$

B Japan's Inheritance Tax — Schedule and Computation

The Japanese government levies an inheritance tax primarily on *heirs* (not on the estate). Tax liability is computed heir-by-heir, using a progressive schedule applied to a notional division of the taxable estate.

Basic Exemption Let H denote the number of statutory heirs (legal heirs under the Civil Code). The basic exemption deducted from the net estate (assets minus debts and funeral costs) is:

$$\text{Basic Exemption} \;=\; 30{,}000{,}000 \text{ JPY} \;+\; 6{,}000{,}000 \text{ JPY} \times H.$$

If the net estate does not exceed this amount, no inheritance tax is due. In addition, there are deductions described in the statutory tax schedule given below.[16]

[16] Tax corresponding to the spouse's statutory share, or up to 160 million JPY, whichever is larger, can be deducted. There are additional per-year amounts for heirs under 18 or with qualifying disabilities. Residents are generally taxed on worldwide assets.

Table A1: Japan Inheritance Tax Schedule (statutory, post-2015)

Taxable Share Bracket (JPY)	**Marginal Rate**	**Deduction (JPY)**
Up to 10,000,000	10%	0
Over 10,000,000 to 30,000,000	15%	500,000
Over 30,000,000 to 50,000,000	20%	2,000,000
Over 50,000,000 to 100,000,000	30%	7,000,000
Over 100,000,000 to 200,000,000	40%	17,000,000
Over 200,000,000 to 300,000,000	45%	27,000,000
Over 300,000,000 to 600,000,000	50%	42,000,000
Over 600,000,000	55%	72,000,000

The distribution of taxable estates Table A2 reports the 2023 distribution of taxable estates by taxable amount, from which we read the two percentiles used in Section 4. The 90th percentile of taxable estates lies between ¥200 and ¥300 million, at approximately ¥250 million. Because the brackets in Table A1 apply to each heir's share rather than to the whole estate, such an estate, split equally among 2.4 heirs after the basic deduction, gives shares of roughly ¥86 million, in the 30 percent bracket. The table also shows what the exemption reform of Section 6 does to the base: cumulative coverage reaches 93.2% at ¥300 million, 97.2% at ¥500 million and 98.5% at ¥700 million, so an exemption of ¥575 million leaves roughly one taxable estate in forty, or about two tenths of one percent of all decedents, subject to the tax.

Table A2: Distribution of taxable estates in Japan (2023)

Taxable amount class (million yen)	Number of decedents	Cumulative count	Cumulative percentage
$\leq$ 50	15,794	15,794	10.1%
50–100	78,684	94,478	60.7%
100–200	40,190	134,668	86.5%
200–300	10,429	145,097	93.2%
300–500	6,248	151,345	97.2%
500–700	1,993	153,338	98.5%
700–1,000	1,210	154,548	99.2%
1,000–2,000	887	155,435	99.8%
2,000–3,000	169	155,604	99.9%
3,000–5,000	81	155,685	100.0%
5,000–7,000	22	155,707	100.0%
7,000–10,000	13	155,720	100.0%
$\geq$ 10,000	20	155,740	100.0%
Total	155,740		

Source: National Tax Agency (2025), Statistics of Taxation, FY2023 data.

Table A2 covers *taxable* estates only. In 2023 there were approximately 1.58 million deaths in Japan, of which 155,740 left taxable estates, so the table describes the top $\approx 9.9\%$ of the overall bequest distribution. The 90th percentile of the overall distribution therefore lies at the taxability margin itself: an estate is taxable when its value exceeds the basic deduction of 30 million yen plus 6 million yen per statutory heir, and with an average of roughly 2.4 statutory heirs per decedent the margin is $30 + 6 \times 2.4 \approx 44$ million yen. Because the margin varies with the number of heirs, from 36 million yen with one heir to 48 million yen with three, we set P_{90} at the round figure of approximately 40 million yen: $40/5 = 8$ times average income ($\bar{y}$, taken as 5 million yen) at the estate level, or, per statutory heir, $40/2.4 \approx 16.7$ million yen ≈ 3.3 times average income—the concept comparable to the De Nardi–Yang target of 4.34, since their Hurd–Smith moment is the estate of a single decedent passing to the model household's one descendant.

C Calibration Details

C.1 The Earnings Process

Our calibration of the labor earnings process is based on the observation that household surveys and administrative tax data provide excellent information on the earnings dynamics of much of the population but not of the richest households (see, for instance, Bosworth and Anders (2011)). To match the earnings dynamics of the whole population, including the very top, we model our earnings process as follows.

First, we approximate the AR(1) process on earnings by a 4-state, first-order Markov process, discretizing the estimated persistence and variance over the four grid points using Tauchen (1986) to make sure that our process accurately represents the estimated earnings dynamics for much of the population. This gives a four-by-four transition matrix.

Second, we add a super-high income state as a fifth income level and assume that only workers in the fourth income state can transit into the super-high income state. We further assume that a super-high income earner can transit into the lowest income state or the fourth income state, or stay in the super-high income state. This leaves three transition parameters, in addition to the level of super-high income. We choose these four parameters by minimizing the mean squared percentage deviation between seven model moments and their data counterparts: the earnings Gini and the shares of total earnings received by the top 1%, 5%, 10%, 20%, 40% and 60%. The model moments are computed over the full population, including retirees.[17] The calibration implies that the super-high state is transitory: a worker who reaches it stays for one year and then returns mostly to the lowest earnings state, so the process captures the very high but short-lived earnings at the top rather than persistent top incomes.

The grid for earnings is given by $\{0.40, 0.74, 1.35, 2.47, 1332.46\}$. The transition matrix for Q_y is:

[17] We also impose adding-up restrictions.

$$\begin{bmatrix} 0.7291 & 0.2664 & 0.0045 & 0.0000 & 0.00000000 \\ 0.1292 & 0.6786 & 0.1901 & 0.0021 & 0.00000000 \\ 0.0021 & 0.1901 & 0.6786 & 0.1292 & 0.00000000 \\ 0.0000 & 0.0045 & 0.2664 & 0.7290 & 0.00013264 \\ 0.9773 & 0.0000 & 0.0000 & 0.0227 & 0.00000000 \end{bmatrix}.$$

The transition matrix for Q_{yh} is

$$\begin{bmatrix} 0.3082 & 0.3801 & 0.2427 & 0.0691 & 0.00000000 \\ 0.2038 & 0.3612 & 0.3111 & 0.1239 & 0.00000000 \\ 0.1239 & 0.3111 & 0.3612 & 0.2038 & 0.00000000 \\ 0.0691 & 0.2427 & 0.3801 & 0.3082 & 0.00000000 \\ 0.0000 & 0.0000 & 0.0018 & 0.9953 & 0.00293748 \end{bmatrix}.$$

The transition matrices induce an initial distribution of earnings with probability masses over the respective earnings levels, given by
$[17.1971\%, 32.8059\%, 32.8029\%, 17.1941\%, 0.00000636\%]$.

Table A3 compares the cross-sectional distribution of earnings generated by our five-state Markov process with its empirical counterpart, computed by Kitao and Yamada (2025).[18] The fit is very good since the Gini coefficient and the top earnings shares are among the moments targeted when choosing the level of the super-high earnings state and the three free parameters that govern transitions into and out of the super-high earnings state. The model reproduces the earnings Gini almost exactly (0.562 versus 0.552 in the data) and matches the shares of total earnings accruing to the top 1% (6.0% versus 6.1%) and the top 5% (19.9% versus 19.5%). The shares of the top 10% through top 80% are likewise within two percentage points of the data, and the top 40% and above within half a point. The calibrated process closely reproduces cross-sectional earnings inequality, including the top earnings shares. This fit does not independently validate upper-tail earnings persistence: the super-high state is effectively a one-year earnings realization.

		Top Percentiles						
	Gini	Top 1%	Top 5%	Top 10%	Top 20%	Top 40%	Top 60%	Top 80%
Data	0.552	6.1	19.5	32.2	52.5	80.9	97.0	100.0
Model	0.562	6.0	19.9	34.1	53.7	81.2	97.2	100.0

Table A3: Moments of earnings inequality: data versus model.

Notes: Shares of total earnings accruing to the indicated top percentiles. Data are the averages of the 2004 and 2014 values from Kitao and Yamada (2025). The Gini and the top 1, 5, 10, 20, 40 and 60 percent shares are targeted in the calibration of the earnings process; the top 80% share is not.

[18] Both the model and data moments are computed using the entire population, including the retirees.

C.2 The Estimated Bequest Function

The two parameters are easier to interpret through the behavior they generate than through their own magnitudes, and following De Nardi (2004) and De Nardi and Yang (2016) we report the two statistics that characterize a luxury-good bequest motive: the level of wealth below which a household leaves nothing, and the marginal propensity to bequeath above it. Both have closed forms in the final period of life, when $p_T = 0$ and the household solves $\max_{a'} U(x-a')+\phi\big(b(a')\big)$ with x denoting cash on hand. The first-order condition is $(b+\phi_2)/c = \kappa$, where

$$\kappa \;\equiv\; \big[(\gamma-1)(-\phi_1)\big]^{1/\gamma} \;=\; \sqrt{43.40} \;=\; 6.59.$$

For estates below the exemption, where $b = a'$, this gives $a' = (\kappa x - \phi_2)/(1+\kappa)$, so bequests begin at a threshold of $\phi_2/\kappa = 0.36$ times average income, or ¥1.8 million, and above that threshold the marginal propensity to bequeath is $\kappa/(1+\kappa) = 0.87$. A household in its last year of life therefore leaves nothing at all if its resources fall below about a third of average income, and devotes 87 of every 100 additional yen above that to its estate.

Above the exemption the marginal yen of estate raises the net bequest only by $1-\tau_b$, so the relevant constant becomes $\kappa_\tau = [(\gamma-1)(-\phi_1)(1-\tau_b)]^{1/\gamma} = 5.12$ and the marginal propensity to bequeath out of cash on hand is $\kappa_\tau/[(1-\tau_b)+\kappa_\tau] = 0.89$ measured against the gross estate, or 0.54 measured against what the heirs actually receive.

The kink at the exemption produces a region of bunching. Setting $a' = x_b$ gives a cash-on-hand level of $[x_b(1+\kappa)+\phi_2]/\kappa = 8.4$ times average income, or ¥41.9 million, at which the optimal estate first reaches the exemption; but the estate remains exactly at x_b until cash on hand reaches $x_b+(x_b+\phi_2)/\kappa_\tau = 8.8$ times average income, or ¥43.9 million. Over that ¥2.0 million window the household leaves precisely the exempt amount and pays no inheritance tax. This is a property of the estimated bequest function together with the estimated exemption, and it is not something the calibration imposes.

D Additional Model Results

D.1 Tax Incidence in the Benchmark Economy

To evaluate the distribution of the tax burden, and to understand how the reforms change it, Tables A4 and A5 report, for selected groups of the wealth distribution, the average age of the group, the share of each tax's total revenue it pays, and the average amount of each tax it pays as a multiple of average income. Age matters here because it is the main determinant of both earnings and accumulated wealth, so it helps explain the pattern of burdens.

Three features of the benchmark economy are worth recording.

First, the inheritance tax is by far the most concentrated of the three instruments. The top 1% of the wealth distribution pays 69.7% of inheritance tax revenue, against

Wealth percentile	Age	Capital tax	Labor tax	Inheritance tax	Total tax
Top 1%	64.68	17.40	0.73	69.72	4.82
Top 1–5%	59.20	12.14	4.71	10.08	5.95
Top 5–10%	58.69	12.49	5.49	10.48	6.66

Table A4: Percentage of the total revenue from a given tax paid by selected wealth percentiles, benchmark economy.

Notes: Each entry is the share of that tax's total revenue paid by the group, in percent. "Age" is the average age of the group.

Wealth percentile	Age	Capital tax	Labor tax	Inheritance tax	Total tax
Top 1%	64.68	0.73	0.17	0.445	1.34
Top 1–5%	59.20	0.13	0.27	0.016	0.41
Top 5–10%	58.69	0.10	0.25	0.013	0.37
All households	52.18	0.04	0.23	0.006	0.28

Table A5: Average amount of a given tax paid by a group, as a multiple of average income, benchmark economy.

Notes: Each entry is the average annual amount of that tax paid by a member of the group, divided by average income in the economy.

17.4% of capital income tax revenue and 0.7% of labor income tax revenue. The next four percent of households pay only 10.1% of inheritance tax revenue, both because they are younger and so less likely to die in a given year and because they are more likely to leave an estate below the exemption.

Second, the wealthy in this economy are old. The average age in the top 1% of the wealth distribution is 64.7, against 52.2 in the population, and most of that group is retired. This is why their share of labor income tax revenue is so small.

Third, the total burden is concentrated but not extremely so. The top 1% pays 1.34 times average income in all taxes each year, against 0.28 for the average household, a ratio of about 4.8. The composition differs sharply across groups: for the top 1% the inheritance tax is the second-largest item after the capital income tax, while for every other group it is negligible.

The second of these points matters for the financing experiments of Section 6.3. If the inheritance tax were reduced and the labor income tax raised to replace the revenue, almost none of the increase would be paid by the households that had been paying the inheritance tax: the two bases barely overlap. This is not a foregone conclusion in a model like ours, since households can become wealthy through very large labor earnings as in Castañeda, Díaz-Giménez, and Ríos-Rull (2003), and such households would face the higher labor tax rate. But in our Japanese calibration the wealthy are overwhelmingly old.

Table A6 reports the same burdens under the reform that raises the exemption to ¥575 million, under both financing instruments.

Wealth percentile	Age	Capital tax	Labor tax	Inheritance tax	Total tax
Capital income tax adjusts					
Top 1%	58.20	0.81	0.22	0.202	1.23
Top 1–5%	60.63	0.15	0.25	0.000	0.39
Top 5–10%	59.18	0.12	0.25	0.000	0.37
All households	52.18	0.05	0.23	0.002	0.28
Labor income tax adjusts					
Top 1%	60.64	0.80	0.20	0.228	1.23
Top 1–5%	60.07	0.14	0.23	0.000	0.37
Top 5–10%	60.04	0.12	0.25	0.000	0.37
All households	52.18	0.05	0.23	0.002	0.28

Table A6: Average amount of a given tax paid by a group, as a multiple of average income, under the reform that raises the exemption to the current U.S. level.

Notes: As Table A5, for the economy in which the inheritance tax rate is held at 39.5% and the exemption is raised to ¥575 million. Small open economy.

Comparing Table A6 with Table A5, the inheritance tax paid by the top 1% falls by roughly half, from 0.445 to between 0.202 and 0.228 times average income depending on the financing instrument, and it falls to zero for every group below the top 1%. The total burden on the top 1% falls only from 1.34 to 1.23 times average income, however, because their capital income tax payments rise with their larger asset holdings. For every other group the total burden is essentially unchanged. The reform is thus narrowly targeted in its incidence: it is a reduction in the tax paid by the wealthiest one percent, largely offset for them by a higher capital income tax bill, and close to invisible for everyone else.

D.2 The Value of Parental Background Under the Exemption Reform

Tables A7 and A8 report the full parental-background matrices behind Section 6.5, for the small open economy with the labor income tax adjusting.

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	0.37	0.27	—	—
4th	0.68	0.57	0.27	—
5th	13.94	13.80	13.35	12.92
Reform, net bequests				
2nd	0.07	—	—	—
3rd	0.38	0.22	—	—
4th	0.77	0.58	0.33	—
5th	13.63	13.39	13.02	12.43
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	0.35	0.21	—	—
4th	0.69	0.54	0.29	—
5th	13.91	13.70	13.34	12.86
Reform, gross bequests				
2nd	0.07	—	—	—
3rd	0.39	0.23	—	—
4th	0.79	0.61	0.34	—
5th	13.85	13.62	13.23	12.65

Table A7: The value of parental background, bequests only: benchmark versus the higher exemption, labor income tax adjusting.

Notes: Each entry is the one-time asset compensation, as a multiple of average income, required to move a newborn from the row background to the column background, switching expected bequests only and keeping the child's initial productivity fixed. The reform holds the inheritance tax rate at its benchmark value of 39.5% and raises the exemption to ¥575 million; the labor income tax adjusts to balance the budget. Small open economy.

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	1.98	1.85	—	—
4th	2.33	2.21	0.27	—
5th	18.03	17.91	15.52	14.88
Reform, net bequests				
2nd	0.07	—	—	—
3rd	1.99	1.78	—	—
4th	2.42	2.21	0.33	—
5th	17.68	17.49	14.94	14.42
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	1.97	1.79	—	—
4th	2.35	2.18	0.29	—
5th	17.93	17.75	15.40	14.80
Reform, gross bequests				
2nd	0.07	—	—	—
3rd	2.01	1.80	—	—
4th	2.46	2.24	0.34	—
5th	17.81	17.62	15.10	14.59

Table A8: The value of parental background, bequests and human capital: benchmark versus the higher exemption, labor income tax adjusting.

Notes: As Table A7, except that each entry switches both the expected bequest and the child's initial productivity.

D.3 Policy Results When Changing the Capital Income Tax

Changing the tax on capital income changes the incentives to save by affecting the net return to accumulation. Changing labor income taxation does not distort labor supply decisions because, for tractability, we assume exogenous labor supply; it does affect net lifetime income, and hence the importance of human capital inheritance across generations.

Aggregate effects. Table A9 reports the aggregate effects in the small open economy when the capital income tax absorbs the change in inheritance tax revenue. The discussion is in Section 6.3; we record the numbers here together with the rate changes required.

	τ_k	W/Y	ΔW	B/W (%)	r_p (%)
Benchmark ($\tau_b = 39.5\%$)	0.220	3.940	—	1.003	3.759
Eliminate ($\tau_b = 0$)	0.249	3.988	+1.2%	0.952	3.779
Cut the rate ($\tau_b = 20\%$)	0.234	3.965	+0.6%	0.975	3.770
Raise exemption (¥575M)	0.216	4.260	+8.1%	1.047	3.884

Table A9: Aggregate effects of inheritance tax reform in the small open economy, adjusting the capital income tax.

Notes: τ_k is the capital income tax rate that balances the government budget. As in Table 5, r_k, K/Y and Y are fixed, so the counterpart of the change in net worth is a change in the net international investment position; r_p is the portfolio return of (13).

The entry worth noting is the exemption reform, which is the only one of the three under which the capital income tax *falls*. The additional accumulation it induces expands the capital income tax base by more than the inheritance tax revenue forgone, so the reform is self-financing at this margin and the lower tax on saving reinforces its direct effect. The consequence is that the exemption reform is the largest of the three under capital-tax financing and the smallest under labor-tax financing.

Distributional effects. Table A10 reports the distributional effects under capital-tax financing. The pattern mirrors the main text, but the movements are larger for the rate reforms: eliminating the tax raises the top 1% share by 2.4 percentage points against 1.3 under labor-tax financing, and raises the Gini from 0.61 to 0.63. The reason is that the higher capital income tax required to replace the revenue bears relatively more heavily on moderate savers in the upper-middle of the distribution than on the super-rich, whose saving is sustained by the bequest motive, so the same reform concentrates wealth more. The exemption reform again lowers top-end concentration, and by slightly more than under labor-tax financing (-0.9 against -0.7 percentage points), because it is the one reform under which the capital income tax falls.

The value of parental background. Tables A11 and A12 recompute the asset compensations of Section 5 for the exemption reform under capital-tax financing, alongside the gross-bequest specification. The pattern is the one described in Section 6.5: the

	Gini	Top 1%	Top 5%	Top 10%	Top 20%
Benchmark ($\tau_b = 39.5\%$)	0.61	17.4	29.3	41.7	60.3
Eliminate ($\tau_b = 0$)	0.63	19.8 (+2.4)	31.8	44.3	62.6
Cut the rate ($\tau_b = 20\%$)	0.61	18.2 (+0.8)	30.2	42.7	61.3
Raise exemption (¥575M)	0.60	16.5 (−0.9)	28.3	41.0	59.9

Table A10: Distributional effects of inheritance tax reform in the small open economy, adjusting the capital income tax.

Notes: As Table 7, with the capital income tax rather than the labor income tax balancing the government budget.

compensations for a super-high background fall slightly relative to the benchmark, from 13.94 to 13.58 times average income on the bequests-only measure and from 18.03 to 17.61 on the total, while the compensations among the bottom four backgrounds rise by a few hundredths. The declines at the top are marginally larger than under labor-tax financing (13.63 and 17.68), which is consistent with the slightly larger fall in the top 1% wealth share in Table A10.

The one place where the sign differs is the gross-bequest specification under capital-tax financing, the last panel of each table: there the compensations for a super-high background *rise*, from 13.91 to 14.68 and from 17.93 to 18.70. This is the same cell of the parameter space in which the welfare verdict of Table 12 turns negative, and for the same reason: the gross-bequest calibration has a higher effective rate, so the reform forgoes more revenue, the required capital income tax increase is larger, and the burden of that increase falls on ordinary savers rather than on the dynastically wealthy.

Parental background in the closed economy. Tables A13 and A14 report the full parental-background matrices for the closed economy, under both financing instruments and both bequest specifications. These are the matrices behind the closed-economy columns of Table 10. Beyond the 5th-state row discussed there, the pattern below the top is the same as in the small open economy: the compensations among the bottom four backgrounds rise by a few hundredths of average income, because the newly exempted estates are those of 3rd- and 4th-state parents.

The division of wealth at death. Table A15 reports how the wealth at death of decedents in the top 1% of the wealth distribution is divided between the net bequest the heirs receive and the inheritance tax.

In the benchmark economy the inheritance tax claims 29.8% of the wealth at death of the richest 1% and the heirs receive 70.2%. The entry worth noting concerns the exemption reform: raising the exemption to the current U.S. level cuts the government's share from 29.8% to about 6.2% without touching the marginal rate, because a ¥575 million exemption shields the bulk of even very large estates. Cutting the rate to 20% leaves the government with about 15%, so the exemption reform relieves these estates more than halving the rate does, even though it leaves the rate alone.

The two financing columns are close but not identical: they agree to within four

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	0.37	0.27	—	—
4th	0.68	0.57	0.27	—
5th	13.94	13.80	13.35	12.92
Reform, net bequests, capital income tax adjusts				
2nd	0.08	—	—	—
3rd	0.39	0.21	—	—
4th	0.77	0.56	0.32	—
5th	13.58	13.31	12.96	12.33
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	0.35	0.21	—	—
4th	0.69	0.54	0.29	—
5th	13.91	13.70	13.34	12.86
Reform, gross bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	0.34	0.20	—	—
4th	0.83	0.66	0.42	—
5th	14.68	14.40	14.09	13.38

Table A11: The value of parental background, bequests only: benchmark versus the higher exemption, capital income tax adjusting.

Notes: As Table A7, with the capital income tax rather than the labor income tax balancing the government budget. The third panel is the gross-bequest economy before the reform, which has its own calibration and so its own benchmark.

hundredths of a percentage point for the benchmark and the exemption reform, and differ by 0.37 points for the rate cut. The incidence of the inheritance tax on the estates it touches is thus governed very largely by the schedule rather than by what happens elsewhere in the budget.

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	1.98	1.85	—	—
4th	2.33	2.21	0.27	—
5th	18.03	17.91	15.52	14.88
Reform, net bequests, capital income tax adjusts				
2nd	0.08	—	—	—
3rd	2.00	1.77	—	—
4th	2.41	2.18	0.32	—
5th	17.61	17.41	14.87	14.36
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	1.97	1.79	—	—
4th	2.35	2.18	0.29	—
5th	17.93	17.75	15.40	14.80
Reform, gross bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	1.97	1.78	—	—
4th	2.52	2.32	0.42	—
5th	18.70	18.53	16.64	15.69

Table A12: The value of parental background, bequests and human capital: benchmark versus the higher exemption, capital income tax adjusting.

Notes: As Table A11, except that each entry switches both the expected bequest and the child's initial productivity.

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	0.37	0.27	—	—
4th	0.68	0.57	0.27	—
5th	13.94	13.80	13.35	12.92
Reform, net bequests, labor income tax adjusts				
2nd	0.08	—	—	—
3rd	0.43	0.26	—	—
4th	0.83	0.62	0.32	—
5th	14.15	13.90	13.46	12.94
Reform, net bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	0.37	0.21	—	—
4th	0.82	0.64	0.39	—
5th	14.18	13.95	13.59	12.97
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	0.35	0.21	—	—
4th	0.69	0.54	0.29	—
5th	13.91	13.70	13.34	12.86
Reform, gross bequests, labor income tax adjusts				
2nd	0.03	—	—	—
3rd	0.35	0.28	—	—
4th	0.76	0.67	0.35	—
5th	14.34	14.24	13.79	13.21
Reform, gross bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	0.43	0.27	—	—
4th	0.85	0.67	0.35	—
5th	14.54	14.33	13.89	13.30

Table A13: The value of parental background in the closed economy, bequests only.

Notes: As Table A7, for the closed economy. Each entry switches the expected bequest while holding the child's initial productivity fixed. The gross-bequest economy has its own calibration and so its own benchmark, reported separately.

	Moving to parent's earnings			
Parent's earnings	1st	2nd	3rd	4th
Benchmark				
2nd	0.04	—	—	—
3rd	1.98	1.85	—	—
4th	2.33	2.21	0.27	—
5th	18.03	17.91	15.52	14.88
Reform, net bequests, labor income tax adjusts				
2nd	0.08	—	—	—
3rd	2.07	1.84	—	—
4th	2.49	2.26	0.32	—
5th	18.25	18.04	15.82	14.91
Reform, net bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	1.98	1.78	—	—
4th	2.49	2.27	0.39	—
5th	18.29	18.09	16.12	14.94
Gross bequest benchmark				
2nd	0.06	—	—	—
3rd	1.97	1.79	—	—
4th	2.35	2.18	0.29	—
5th	17.93	17.75	15.40	14.80
Reform, gross bequests, labor income tax adjusts				
2nd	0.03	—	—	—
3rd	1.97	1.87	—	—
4th	2.42	2.32	0.35	—
5th	18.41	18.32	16.37	15.21
Reform, gross bequests, capital income tax adjusts				
2nd	0.07	—	—	—
3rd	2.07	1.87	—	—
4th	2.53	2.32	0.35	—
5th	18.62	18.44	16.46	15.46

Table A14: The value of parental background in the closed economy, bequests and human capital.

Notes: As Table A13, except that each entry switches both the expected bequest and the child's initial productivity.

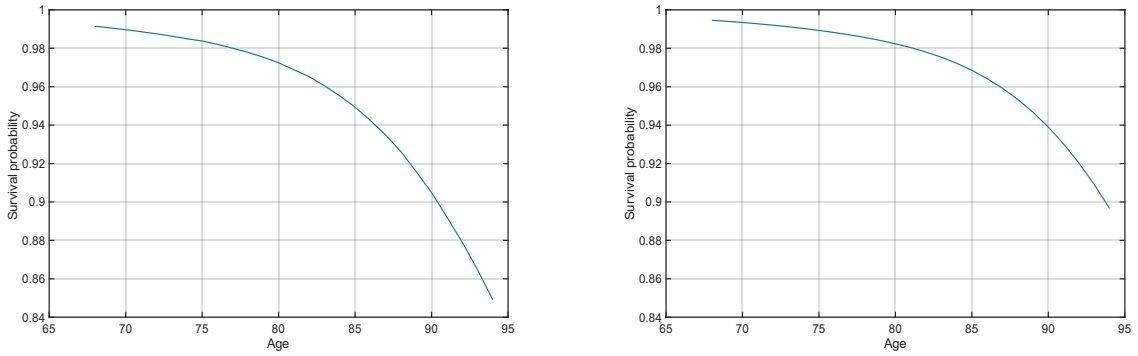
	Capital income tax adjusts		Labor income tax adjusts	
	Heirs receive	Tax claims	Heirs receive	Tax claims
Benchmark ($\tau_b = 39.5\%$)	70.22%	29.78%	70.22%	29.78%
Cut the rate ($\tau_b = 20\%$)	85.27%	14.73%	84.90%	15.10%
Raise exemption (¥575M)	93.86%	6.14%	93.83%	6.17%
Eliminate ($\tau_b = 0$)	100.00%	0.00%	100.00%	0.00%

Table A15: Division of the wealth at death of the top 1% between net bequests and inheritance taxes.

Notes: Shares of the wealth at death of decedents in the top 1% of the wealth distribution. Small open economy.

D.4 Survival Probabilities in the Population-Aging Experiment

Figure A2 compares the conditional survival probabilities used in the benchmark with those used in the population-aging experiment of Section 10.



(a) Benchmark: 1990 birth cohort (b) Low mortality: 2070 birth cohort

Figure A2: Conditional survival probabilities by age in the benchmark and in the population-aging experiment.