

When Interest Rates Go Low, Should Public Debt Go High?*

Johannes Brumm
Karlsruhe Institute of Technology
johannes.brumm@gmail.com

Xiangyu Feng
Xiamen University
xyfeng@bu.edu

Laurence Kotlikoff
Boston University
kotlikoff@gmail.com

Felix Kubler
University of Zurich
fkubler@gmail.com

October 17, 2022

Abstract

Is deficit finance, explicit or implicit, free when borrowing rates are routinely lower than growth rates? Specifically, can the government make all generations better off by perpetually taking from the young and giving to the old? We study this question in stochastic OLG models oriented toward this outcome. Unfortunately, Pareto gains are predicted only for implausible calibrations. Even then, the gains reflect improved intergenerational risk-sharing, improved international risk-sharing, and beggaring thy neighbor – not intergenerational redistribution per se. Low government borrowing rates justify state-contingent transfers between generations, not unconditional redistribution from young and future generations.

Keywords: dynamic efficiency, fiscal policy, government debt, real interest rate, risk premium, risk-sharing, social security.

JEL Classification Codes: E43, E62, H63.

*We thank Oliver Blanchard, Werner Roeger, and seminar participants at the Bank of International Settlements, the European Central Bank, the European Commission, Goethe University Frankfurt, Harvard University, the International Monetary Fund, and the Society for Economic Dynamics 2021 meeting for very helpful comments. Johannes Brumm acknowledges financial support from the ERC. Xiangyu Feng acknowledges financial support from the Fundamental Research Funds for the Central Universities (2072021144).

1 Introduction

When interest rates go low, should public debt go high? Given the positive average gap between the U.S. growth rate and the real interest rate on U.S. Treasuries, Blanchard (2019) suggests that deficit finance, explicit or implicit, a) has, on average, no fiscal cost and b) entails negligible and, potentially, negative welfare costs. The first claim has received qualified confirmation.¹ The second claim, studied here, has received less attention.² We revisit Blanchard’s simple closed-economy OLG model and also examine an open-economy variant with three questions in mind. First, can deficit finance effect Pareto improvements in ex-ante expected utility levels of current and future generations? Second, do potential Pareto gains reflect intergenerational redistribution per se or improved intergenerational risk-sharing? Third, is there more scope for domestic Pareto gains in open economies, and, if so, is this due to beggaring thy neighbor’s capital stock or improving international risk-sharing?

To preview our findings, deficit finance, which we implement via a defined-benefit, pay-go, social security system (henceforth, simply, pay-go) can Pareto improve Blanchard’s economy, but only under, in our view, implausible assumptions. Moreover, when such efficiency gains arise, improved risk-sharing is the reason, reducing both the risk of being born under adverse conditions and retiring under such conditions. Indeed, in Blanchard’s closed-economy model, if one first implements a Pareto-improving, risk-sharing policy – in which the young give to the old in some states and receive from the old in other states – adding his proposed deficit policy to the mix serves only to redistribute, helping early generations at the price of hurting future generations.

In the open economy, deficit finance has more potential to generate domestic welfare gains, albeit at a potential net cost to foreigners. Open economies feature globally mobile capital. Hence, domestic deficits crowd out the global capital stock. Domestic “Pareto gains” from running deficits may thus largely represent beggaring-thy-neighbor, i.e., redistributing from foreigners by permitting the home country to consume output, most of which would otherwise be invested abroad.³ On the other hand, we show that the presence of foreign economies expands the scope of risk sharing even if only domestic transfer schemes are considered.

We make three key points by recapitulating the simple case of linear technology under which no crowding out occurs. First, absent crowding out, if the growth rate is certain and exceeds the risk-free rate, deficit finance is Pareto improving, i.e., non-state-contingent redistribution from the young to the old makes all generations better off. Second, state-contingent transfer policy Pareto dominates deficit finance independent of the risk-free rate. Third, Pareto improvements need not benefit particular groups. Yes, efficiency gains can be allocated largely to initial generations, but they can also be spent on children and future generations. Points two and three remain true in a model that allows for crowding out, for which we characterize Pareto

¹See, e.g., Reis (2021), Mian et al. (2021a), and references therein.

²Section 2 reviews some of the large literature on dynamic efficiency and risk-sharing in OLG models on which Blanchard (2019) and this paper build.

³Note that assuming linear technology, the first case we consider and Blanchard (2019)’s focal case, implicitly treats the domestic economy as tiny and open. It thus maximizes beggar-thy-neighbor.

efficient allocations and the associated state-contingent transfers in the constant relative risk aversion (CRRA) case – discussed in the text and derived in the Appendix. Thus, on theoretical grounds, pay-go is not the natural response to inefficiencies suggested by low safe rates. The rest of the paper studies whether this conclusion still holds in realistically calibrated closed and open economies with recursive utility.

We begin by modifying Blanchard (2019)’s version of Epstein and Zin (1989) preferences to ensure that the degree of risk aversion concerning one’s birth state coincides with the degree of risk aversion concerning one’s old-age consumption. This influences neither agents’ behavior nor the model’s calibration. But it does provide greater scope for Pareto improvements via deficit finance, which makes our results even more striking.

Next we use his closed economy to show the importance of risk-sharing to the success of pay-go. We do so in three ways. First, we decompose a generation’s ex-ante expected utility (henceforth EAU) gains into three components – a risk-neutral effect, a life-cycle-risk effect, and a cohort-risk effect. The risk-neutral effect (RNE) isolates the EAU impact of policy-induced changes in average, realized, cohort-specific consumption levels. The life-cycle-risk effect (LRE) captures the EAU impact of old-age consumption risk given an agent’s state of birth. And the cohort-risk effect (CRE) captures the EAU impact of the risk faced by cohorts of being born in a bad state. Our decomposition relates the percent change in EAU to the sum of RNE, LRE, and CRE. As we show, increases in EAU of those alive in the long run (henceforth, long-run EAU) are fully explained by LRE and CRE, which capture improved risk-sharing both post- and pre-birth.⁴

Second, we consider a two-way transfer scheme where the old transfer to the young when shocks are good and the young transfer to the old when shocks are bad – with the policy redistributing nothing, on average, across generations. We show that such a scheme a) is Pareto improving when pay-go is not, b) two-way transfers Pareto dominate pay-go even when pay-go itself is Pareto improving, and, once implemented, c) exhausts the potential for additional Pareto gains via pay-go. Finally, as our third means of showing that risk-sharing underlies successful pay-go, we show that Pareto gains, if they arise, disappear in the absence of Blanchard (2019)’s assumption that the young receive a safe endowment in addition to their risky wages.⁵ Clearly, dropping the assumption of a fixed-wage component limits the ability of the young to share risk with the old.

In short, in Blanchard (2019)’s model, the source of potential Pareto improvements, under the extreme conditions that they arise, is intergenerational risk sharing. An average risk-free rate that is low signals the presence of risk – risk that can potentially be shared across generations, with state-contingent transfers the clear policy response.

Our open-economy model takes the U.S. as the home country and the RoW as the foreign country. The model includes compensated foreign-investment taxes to achieve a realistic pattern

⁴Since increases in EAU decline monotonically as one moves from early to later generations, increases in long-run EAU ensures that all prior generations are also better off, i.e., it implies a Pareto improvement in the path of EAU.

⁵By analogy, Samuelson (1958)’s Pareto improvement from his Ponzi scheme disappears if one drops his no-storage assumption.

of cross-country asset holdings.⁶ Opening up the economy expands the scope for domestic efficiency gains. That is, generating higher EAU levels for domestic residents is feasible for a somewhat larger range of parameters than in the closed economy. However, this home-country improvement typically reflects the aforementioned beggar-thy-neighbor effect. Indeed, the U.S. experiences only about one quarter of the reduction in domestic capital that would arise were the economy closed. Stated differently, the RoW takes most of the hit to global investment. But given our assumed incomplete, international financial markets, domestic pay-go also benefits foreigners through improved international risk-sharing. Specifically, the reduction in saving in the home country reduces the global demand for the domestically-supplied safe bond, raising its return. This lets foreign agents build safer portfolios at lower cost. With sufficiently low risk-free rates as well as exceedingly small risk premia, both countries can gain from domestic deficits. Yet, here again, state-contingent two-way transfers Pareto-dominate pay-go, delivering substantial risk-sharing benefits with much less crowding out, for both domestic and foreign agents.

To summarize, in the special cases in which they occur, the secret sauce underlying domestic Pareto-improving deficits in Blanchard’s model is improved risk-sharing, intergenerationally and internationally, plus begging thy neighbor. An appropriate risk-sharing policy can make all generations, foreign and domestic, young and old, better off.

Section 2 proceeds with a brief literature review. Section 3 provides simple examples to clarify key elements underlying our analysis. Section 4 presents the closed-economy model, including its calibration and pay-go impacts. Section 5 decomposes EAU changes into risk neutral and risk-sharing components and demonstrates the role of risk-sharing in several ways. Section 6 adds the rest of the world as a second country. Section 7 acknowledges that deficit finance may be justified in Keynesian and other settings even if it isn’t justified in Blanchard (2019)’s. Section 8 summarizes and concludes.

2 Literature Review

This paper belongs to the large literature on efficiency in OLG models, both theoretical and quantitative, and to the recent literature on government debt in low interest rate environments.⁷ As Samuelson (1958) and Diamond (1965) showed, OLG economies aren’t necessarily Pareto-efficient. In deterministic OLG models, pay go is the clear policy response to “dynamic inefficiency.” Add uncertainty and matters become far more complex, starting with two competing notions of Pareto-efficiency – interim efficiency and ex ante efficiency. Interim efficiency references making agents better off given the state of nature those alive are currently in and the states of nature those not yet born find themselves when they are born⁸. Making agents better off means, in the case of newborns, raising their expected lifetime utility and, in the case of

⁶Assuming real transaction costs of investing abroad would serve equally well.

⁷Reis (2022) and Blanchard (2023) provide excellent literature reviews.

⁸Since we consider only adults, state of birth references the economy’s and individual’s economic position upon entering the workforce.

those born in prior periods, raising their expected remaining expected lifetime utility. Ex ante efficiency is identical to ex interim utility for those alive. For the unborn, it references increasing current expected lifetime utility, assessed over all the states of nature into which they may be born.

Whereas ex interim and ex ante expected utility are clearly defined, there are different notions of dynamic efficiency. Abel et al. (1989) use this term to reference interim Pareto efficiency, for which they derive a sufficient condition, namely that output net of wages is higher than investment in all date-events. Hellwig (2021) calls an allocation dynamically efficient if it cannot be improved using non-contingent transfers between young and old, holding investment fixed. Whether this holds depends only on how the risk-free rate compares to the growth rate of the economy. In contrast, Zilcha (1990, 1991) and Barbie et al. (2007) generalize the production-based definition of Cass (1972) to the stochastic case and call a competitive equilibrium dynamically efficient if and only if it is not possible to increase aggregate consumption at some date-event without reducing it at any other date-event. We adopt this definition.⁹ Interestingly, dynamically efficient equilibria of stochastic OLG models with production may fail to be interim Pareto efficient, let alone ex ante Pareto efficient. This holds even if financial markets are (sequentially) complete. Under the ex-ante criterion, competitive equilibria in stochastic OLG models are inherently Pareto inefficient as generations born at different date-events in the future cannot share risk. Gordon and Varian (1988), Ball and Mankiw (2007), and others provide general conditions for ex-ante Pareto-efficient intergenerational risk-sharing, while Gottardi and Kubler (2011) explore the scope for ex-ante Pareto-improving social security using simple examples.

Other relevant studies are quantitative, assessing the potential of pay go to improve the intertemporal allocation. After consulting the historical record, Abel et al. (1989) conclude that the U.S. and all other major OECD economies are interim Pareto efficient. Krueger and Kubler (2006) distinguish social security's separate roles in sharing risk and reducing capital formation. Risk, in their model, arises from macro shocks that can drive wages and capital returns in different directions. They find pay-go social security is interim Pareto improving in partial, but not general equilibrium. Hasanhodzic and Kotlikoff (2013)'s yearly calibrated OLG model shows that intergenerational risk, caused by macro shocks, is far smaller than suggested by models with fewer periods.¹⁰ İmrohoroglu et al. (1995) and Hubbard and Judd (1987) examine the state's ability to share idiosyncratic micro risk, particularly longevity risk. They demonstrate that risk-sharing associated with state pensions can raise long-run welfare despite the policy's crowding out of capital. Harenberg and Ludwig (2019) reach the same conclusion by combining correlated micro and macro shocks. The interaction of these shocks substantially exacerbates aggregate risk, making risk mitigation more important than crowding out in determining the long-term gains from pay-go social security.

⁹Barbie et al. (2007) provide full characterizations in terms of equilibrium prices of both dynamic efficiency and interim Pareto-efficiency. Unfortunately, these conditions are hard to apply as they involve convergence of infinite sums.

¹⁰If, as claimed here, risk-sharing is the sine qua non for Pareto-improving deficit finance, a dearth of risk to share in realistically-timed models raises further doubt about the efficiency of intergenerational redistribution.

Blanchard (2019) ties the question of government debt and intergenerational transfers to the current low interest-rate environment.¹¹ He, like Sergeyev and Mehrotra (2021) and Summers and Rachel (2019), argues that deficits may entail no fiscal costs. Brumm et al. (2022) provides several stylized counterexamples showing that seemingly free deficits may be nothing of the sort. Brumm et al. (2020), Evans (2020) and Hasanhodzic (2020) critically discuss some of the assumptions underlying Blanchard (2019).¹² Barro (2020) traces low risk-free rates to rare disaster risk.¹³ Ball and Mankiw (2021) show, in a deterministic model with market power, that deficit finance may reduce welfare even with very low safe rates. Since risk and market power are certainly two of the major drivers of the wedge between government borrowing rates and the marginal product of capital, their findings complement ours in questioning whether low safe rates justify higher deficit spending. These concerns are relevant even if growth rates in excess of safe rates allow governments to run, on average, large deficits. Other studies examine whether low rates support pay go when government bonds provide liquidity services (see Sims 2019 and Mian et al. 2021a) or help insure either against idiosyncratic risk (see Reis 2021, Miao and Su 2021, and Brunnermeier et al. 2020) or aggregate risk (see Abel and Panageas 2022).

3 Intergenerational Transfers and Risk-Sharing

This section first presents the basic model of our paper, which will be extended to incorporate Epstein-Zin preferences in section 4 and a second country in section 6. Next, we recapitulate the simple case of linear technology under which no crowding out occurs. We show that pay-go policy is Pareto improving if the growth rate exceeds the risk-free rate. Yet, pay-go is always Pareto-dominated by state-contingent and, generically, two-way transfer policy. Finally, we allow for crowding out and discuss Pareto efficient allocations along with the associated state-contingent transfers, which we formally characterize in Appendix A.

3.1 OLG Economy With Intergenerational Transfers

Agents, with mass 1 per generation, live for two periods, supplying labor inelastically when young and consuming when young and old. Without loss of generality, there is neither population nor productivity growth. Aggregate output is given by a fixed endowment, E , received by the young, as in Blanchard (2019), and a Cobb-Douglas production with total factor productivity, $A_t > 0$, drawn from a distribution with finite support:¹⁴

$$f(A_t, k_t) = A_t k_t^\alpha,$$

¹¹Blanchard et al. (2020) explore some implications of these arguments for EU fiscal rules.

¹²None of these papers consider either the importance of risk-sharing or beggar-thy-neighbor policy for assessing deficit policies.

¹³We assume rare disasters in appendix D, which allows us to match the targeted risk premium with a much lower risk aversion while not substantially changing the welfare implications of pay-go policies.

¹⁴Starting from section 4, productivity is distributed log-normally as in Blanchard (2019).

where k_t denotes the time t capital stock and $\alpha \in [0, 1]$. When young, agents receive T , earn wage W_t , and pay a net tax, T_t . When old, agents receive a net transfer, T_{t+1} , and returns to their savings. Each period's transfers are financed via lump-sum taxes on the young. There are two assets – capital, whose return is risky, exceeding the safe return in certain states, and whose principal depreciates fully each period, and a risk-free bond in zero net supply. Agents invest in both assets to maximize time-separable expected utility over consumption when old and consumption when young

$$U(c_y, c_o) = (1 - \beta)v(c_y) + \beta\mathbb{E}[v(c_o)],$$

where $\beta \in (0, 1)$ and $v(\cdot)$ is strictly increasing, strictly concave, and satisfies Inada conditions. The generation born at t solves

$$\begin{aligned} \max_{c_{y,t}, c_{o,t+1}, k_{t+1}, b_{t+1}} \quad & U(c_{y,t}, c_{o,t+1}) \\ \text{s.t.} \quad & c_{y,t} = W_t + E - k_{t+1} - b_{t+1} - T_t \\ & c_{o,t+1} = k_{t+1}R_{t+1} + b_{t+1}R_{t+1}^f + T_{t+1}. \end{aligned} \tag{1}$$

Prices R_t and W_t satisfy

$$R_t = R(A_t, k_t) = \alpha A_t k_t^{\alpha-1}. \tag{2}$$

$$W_t = W(A_t, k_t) = (1 - \alpha)A_t k_t^\alpha. \tag{3}$$

The supply of capital in t equals the savings of the young in $t - 1$. The gross risk-free rate is given by

$$R_{t+1}^f = \frac{(1 - \beta)v'(c_{y,t})}{\beta\mathbb{E}_t[v'(c_{o,t+1})]}. \tag{4}$$

We contrast the laissez-faire economy, $T_t = 0$, with a tax-transfer (pay-go) policy that is constant over time, $T_t = T > 0$, with a tax-transfer policy that can depend on the current (productivity) shock, $T_t = T(A_t)$, and with a tax-transfer policy that depends on the shock and on the beginning-of-period capital stock, $T_t = T(A_t, k_t)$.

Note that we can think of such a tax-transfer policy, $\{T_t\}_{t=0}^\infty$, as a debt policy augmented by a (small) lump-sum tax policy, see Green and Kotikoff (2008). In period t , the government borrows $D_t = T_{t+1}/R_{t+1}^f$ from the young to finance debt repayments to the old, who get $R_t^f D_{t-1} = T_t$ as they had lend $D_{t-1} = T_t/R_t^f$. To balance its budget at t , the government levies (potentially negative) lump-sum taxes $R_t^f D_{t-1} - D_t$ from the young. For $D_t = D$ and thus $T_t = R_t^f D$ these taxes are positive (negative) exactly when the risk-free rate is. This tax-transfer policy is called constant debt policy by Blanchard (2019) and is analyzed in Appendix E.

3.2 Pay-go – a Simple Example

Take the simplest case of $\alpha = 1$, which renders factor prices independent of the capital stock and, thus, abstracts from crowding out. Also assume two realizations of the productivity shock,

$A_H > A_L > 0$, each occurring with probability one half. Take $T_t = T$. Then pay-go is Pareto improving if

$$\frac{dEU}{dT} = -(1 - \beta)v'(c_y) + \frac{\beta}{2}(v'(c_{o,H}) + v'(c_{o,L})) > 0,$$

which holds if the risk-free rate is negative (see equation 4). This unambiguous result hinges on the assumption of linear production, which is unrealistic except as a stand in for assuming the economy is small and open.¹⁵

Note that the gains from pay-go need not be distributed to the initial old to any degree, but can, instead, be used to benefit one or more later generation to smaller or larger extents. In other words, the observation that the growth rate exceeds the safe rate does not militate toward redistribution to initial generations. As one example, the government could start pay-go at $t = n$ rather than $t = 0$. This also constitutes a Pareto improvement. When started at n , the generation born at n receives T in old age without having contributed to the system. This maximizes their pay-go gains relative to other start dates. Indeed, in this stationary setting, if T is set such that $R_t^f = 1$ for all $t \geq n$, any initial start date, n , achieves a constrained Pareto optimum, where constrained references exclusively running pay-go policy. Hence, the ability, when $\alpha = 1$, to achieve a constrained Pareto optimum by running standard pay-go does not require starting the policy immediately.¹⁶

A different issue is whether pay-go is efficient. It's not. Adding an additional shock-dependent transfer always improves the allocation. To see this, note that with a fixed transfer, consumption when young is constant and $v'(c_{o,H}) < v'(c_{o,L})$. Therefore, adding and subtracting an infinitesimal transfer τ to the transfer T in the low and high state, respectively, changes the expected old-age utility of the current young (who will not pay any contingent transfer when young, but receive it when old) by

$$\frac{dEU^O}{d\tau} = \frac{\beta}{2}(-v'(c_{o,H}) + v'(c_{o,L})) > 0,$$

and all future agents' utility by

$$\frac{dEU}{d\tau} = -\frac{(1 - \beta)}{2}v'(c_y)(\tau - \tau) + \frac{dEU^O}{d\tau} = \frac{dEU^O}{d\tau} > 0,$$

independently of the risk-free rate.¹⁷

¹⁵If an open-economy model (with $\alpha < 1$) is explicitly modeled, as we do in section 6, a negative risk-free rate per se no longer implies that pay-go Pareto improves – it also depends on the risky rate of return. Moreover, the policy may “Pareto improve” domestic agents' welfare only at the cost of making foreigners worse off.

¹⁶Alternatively, suppose the government, starting at $t = 0$, taxes T , which is infinitesimally small, from each worker. Starting at $t = 1$ it transfers TR_{t+1}^f to the old. As such, this scheme leaves the welfare of all generations unchanged. However, this scheme frees up resources for the government to invest in capital, namely T at $t = 0$ and $T(1 - R_t^f)$ at all $t > 1$. Thus, the government can, for instance, transfer $TR_1 \dots R_n + T(1 - R_1^f)R_2 \dots R_n + \dots + T(1 - R_{n-1}^f)R_n$ to those born at $t = n$. At $t = n + m$ the government has once again accumulated $T(1 - R_n^f)R_{n+1} \dots R_m + \dots + T(1 - R_{m-1}^f)R_m$ which it could transfer to those born at $t = m -$ and so on. Clearly, as n and m are arbitrary, there is a continuum of Pareto improvements based on this or alternative deterministic or randomized policies.

¹⁷We assume that the introduction of transfers is announced one period ahead, but after consumption takes place.

In words, if consumption when old only depends on the current shock, a shock-dependent transfer scheme can always make everybody better off. This raises the question of whether we can achieve a Pareto-efficient allocation, as opposed to just a Pareto improvement, with only current-shock contingent transfers. The answer is generally no. Given any shock-contingent transfers with a resulting allocation $(c_{y,H}, c_{y,L}, c_{o,HH}, c_{o,HL}, c_{o,LH}, c_{o,LL})$, $c_{o,LH} < c_{o,HH}$ will generally hold because when the previous shock was L the young need to transfer to the old, while with a H shock they receive transfers and hence save more. But then there always exist transfers that depend on the current and the previous shock that make everybody better off. The construction is exactly as above. The same argument holds for any scheme where transfers depend on a finite history of shocks.¹⁸

3.3 Pareto-efficient Transfers

As our simple example illustrates, efficient allocations depend on the entire history of shocks, not a subset. We now turn to the general case where $\alpha < 1$ and thus crowding out occurs. We denote a history of shocks by $A^t = (A_0, \dots, A_t)$ and aggregate consumption at any date event A^t by $C(A^t)$. In Proposition 1, proved in Appendix A, we characterize, for CRRA utility, the set of efficient allocations and their decentralization via tax-transfer policy.

Proposition 1: *Suppose agents' utility exhibits constant relative risk aversion.*

(a) *For any Pareto-efficient allocation and all t , there exists a $\lambda_t \in [0, 1]$, such that*

$$c_y(A^t) = \lambda_t C(A^t), \text{ for all } A^t.$$

(b) *For $E > 0$, no laissez-faire equilibrium is Pareto efficient.*

(c) *For any sequence $(\lambda_t)_{t=0}^{\infty}$ with $\lambda_t \in (0, 1)$ and $\frac{\beta}{1-\beta} \left(\frac{\lambda_t}{1-\lambda_{t+1}} \right)^\gamma < 1$, for all t , there is a Pareto efficient allocation that can be implemented as a competitive equilibrium with lump-sum transfers by a time-varying transfer scheme*

$$T_t(k, A) = (1 - \lambda_t) (E + W(A, k) - S_t(A, k)) - \lambda_t k R(A, k),$$

with the savings functions $S_t(A, k)$ solving

$$\begin{aligned} & (1 - \beta) \lambda_t^{-\gamma} (E + f(A, k) - S_t(A, k))^{-\gamma} \\ & = \beta (1 - \lambda_{t+1})^{-\gamma} \sum_{A'} \pi_{A'} R(A', S_t(A, k)) (E + f(A', S_t(A, k)) - S_{t+1}(A', S_t(A, k)))^{-\gamma}. \end{aligned}$$

We show that efficiency entails the choice of a date-specific, but not state-specific linear

¹⁸This is consistent with the analysis in Gottardi and Kubler (2011) that shows in this framework that an agent's consumption in any efficient allocation at any time t can be expressed as a function of the current shock and aggregate consumption.

sharing rule for aggregate consumption.¹⁹ To be precise, let λ_t reference the share of aggregate consumption consumed by the young at time t . Then the sequence of these shares can be chosen to achieve a Pareto improvement relative to *laissez faire*. Yet, as there is a continuum of inefficient Pareto improvements, there is a continuum of efficient Pareto improvements. Efficiency gains can, for example, be solely allocated to the initial elderly, leaving initial young and future generations at their *laissez-faire* levels of EAU. Alternatively, none of the gains can be allocated to the initial old with initial young and future generations sharing all achievable gains in innumerable different ways. Thus, our formal characterization of the set of Pareto efficient allocations also does not militate toward favoring the initial elderly. What it does show, however, is that the transfers required to reach such allocations are, in general, two-way and highly state dependent even if the risk-free rate is negative. As we show numerically, these insights extend to realistically calibrated models with recursive utility.

4 Blanchard's Closed Economy Revisited

Here we revisit Blanchard (2019)'s closed economy. But we replace, in section 4.1, Blanchard's welfare measure with one that fully accords with the assumed Epstein-Zin preferences. Doing so expands the set of parameter values for which deficit finance is Pareto improving. Even so, this set turns out to be quite narrow, comprising very low risk-free rates and/or risk-premia, as we report in section 4.3.

Blanchard (2019)'s OLG model is the model presented in section 3.1, except for preferences and the specification of productivity shocks. Agents' utility from consuming $c_{y,t}$ when young and $c_{o,t}$ when old is homothetic Kreps-Porteus, with an intertemporal elasticity of substitution (IES) of 1, a risk-aversion parameter denoted γ , and a discount rate of $\tilde{\beta} = \beta/(1-\beta)$. Blanchard (2019) specifies these preferences with the utility function

$$(1 - \beta) \log c_{y,t} + \frac{\beta}{1 - \gamma} \log \mathbb{E}_t [c_{o,t+1}^{1-\gamma}]. \quad (5)$$

By a monotone transformation via the exponential function we get:

$$c_{y,t}^{(1-\beta)} \mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}^{\frac{\beta}{1-\gamma}}. \quad (6)$$

This is in line with Epstein and Zin (1989)'s original formulation. It also has the advantage of being homogeneous of degree one so that variations in utility correspond to variations in consumption, i.e., a given percentage increase in consumption in all states produces the same percentage increase in utility.

Productivity shocks are assumed to be log-normally distributed:

$$\log A_t = \epsilon_t, \quad \epsilon_t \sim_{i.i.d.} \mathcal{N} \{0, \sigma^2\}. \quad (7)$$

¹⁹Ball and Mankiw (2007) suggests this result.

Details on approximating expectations over the productivity shock and computing equilibria can be found in appendix B.

As above we contrast the laissez-faire economy, $T_t = 0$, with a tax-transfer policy that is constant over time, $T_t = T > 0$. We call this policy (defined-benefit) pay-go. While we follow Blanchard (2019) in focusing on this transfer scheme, we also consider other schemes – defined contribution, constant debt, and two-way transfers (see appendix E for the first two and sections 5.2 and 6.3 for the latter).

4.1 Ex-Ante Utility

Pay-go policies redistribute to the initial elderly. Hence, the crucial question is how such policies affect future generations. We assess their welfare using the following ex-ante expected utility function:

$$U_0^t = \left(\mathbb{E}_0 \left[\left(c_{y,t}^{(1-\beta)} \mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}^{\frac{\beta}{(1-\gamma)}} \right)^{1-\gamma} \right] \right)^{\frac{1}{1-\gamma}}, \quad (8)$$

where 0 is the time of assessment, i.e., when a policy choice is made, and $t > 0$ is the time of birth. This measure evaluates uncertainty about the state in which an agent is born with the same degree of risk aversion with which the agent evaluates the uncertainty of old-age consumption. Blanchard (2019), in contrast, evaluates long-run EAU via

$$\tilde{U}_0^t = \mathbb{E}_0 \left[(1 - \beta) \log c_{y,t} + \frac{\beta}{1 - \gamma} \log \mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \} \right] = \mathbb{E}_0 \left[\log \left(c_{y,t}^{(1-\beta)} \mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}^{\frac{\beta}{(1-\gamma)}} \right) \right], \quad (9)$$

which effectively assumes risk aversion of 1 with respect to the state in which an agent is born. That’s much lower than the risk aversion re old-age consumption needed to match the risk premium in the considered calibrations. Because of this difference, our measure (8) is more favorable to finding welfare improvements from deficit finance than is (9), as we illustrate in appendix C.1.²⁰ Appendix C motivates the use of EAU as defined in (8) by ways of a simple example and provides a derivation of (8) from basic assumptions.

4.2 Calibration

Following Blanchard (2019), we set the capital share, α , at 0.33, the fixed endowment, E , at average wages in the no-policy stochastic steady state, and we consider, except in section 5.3, only cases of $T \leq E$, which ensures feasibility.²¹ Also, following Blanchard (2019), we set the standard deviation of the productivity shock, σ , at = 0.2, which he chooses as a compromise between the lower value implied by data on TFP growth volatility and the higher value needed to match the volatility of stock returns.

²⁰But, to repeat, the choice between (8) and (9) has no impact on the calibration, as both specifications reflect the same preferences at the ex-interim stage when agents are alive and making choices.

²¹For $T \geq E$, there will be cases of *game over* – realizations of A_t in which the young have too few resources to cover their pay-go contributions. This “game over” limit, examined by Evans et al. (2012), plays a key role in Tirole (1985) and other studies of bubbles of finite value.

To calibrate preference parameters, we first fix a pair of targets for the unconditional mean of the risk-free rate (RFR), $\mathbb{E}_0 [R_t^f]$, and the risk premium (RP), $\mathbb{E}_0 [R_t - R_t^f]$. We then choose pairs of β and γ that meet the targets. We focus on two cases. Baseline 1 has an annualized risk-free rate of -1 percent and an annualized risk premium of 3 percent. To hit these targets, we need γ to equal 19.0 and an annualized $\tilde{\beta}$ of 0.933. Baseline 2 features an even lower risk-free rate of -2 percent and a risk premium of 3. Here γ is set at 19.2 and $\tilde{\beta}$ at 0.944. The values of γ and β needed to hit these two targets are quite sensitive to TFP risk. Fortunately, as shown in appendix D, our main results aren't particularly sensitive to changes in σ and the associated changes in γ and β .

Blanchard's calibration abstracts from both population and TFP growth. The average postwar U.S. population growth rate was roughly 1 percent and the average growth rate of TFP was around 1.5 percent. Hence, a -2 percent differential between the average safe rate and the average growth rate corresponds to an annual safe rate of about 0.5 percent in a model where population growth and TFP growth are matched to historical averages.²² The historical average real return on the 1-year U.S. Treasury bill rate is 0.6 percent. Hence, calibrating, as we do, a real risk-free rate/growth rate differential of -1 percent and -2 percent (net of growth) in our two baselines (which we call B1 and B2) appears to capture the range of empirically plausible parameters.

What about the risk premium? The historical average risk premium on equity has been well above 4 percent. On the other hand, returns to physical capital as measured from national product accounts seem to lie slightly below 4 percent. Of course, physical capital is just a portion of U.S. national wealth, whose real return has averaged 6.5 percent in the postwar era. It averaged 9.5 percent between 2010 and 2019.²³ In sum, our baseline assumption of a 3 percentage point (pp) risk premium seems at the low end of what's empirically reasonable. Nonetheless, we adopt this value to give deficit finance the benefit of the doubt. For, as we and Blanchard (2019) show, adopting a higher and, to us, more plausible risk premium rules out Pareto-improving deficit (pay-go) policy. Note, though, our analysis does not simply reproduce Blanchard's due to our adoption of internally consistent preferences – preferences, which, as mentioned, are more likely to admit Pareto improvements from pay-go.

These crucial calibration targets complete our description of the closed-economy's calibration. For the open-economy cases, most of the calibration details carry over. Remaining details are described in section 6.

4.3 Role of the Risk-Free Rate, Risk-Premium, and Policy Scale

We now consider the EAU impact of introducing pay-go policy. Following Blanchard (2019), each young cohort pays the old a fixed amount, set at 20 percent of average capital in the no-policy, stochastic steady state. Unless otherwise stated, all results presented below reflect

²²This said, U.S. population growth is far from stationary, and is projected to decline to zero in the second half of this century, see Aksoy et al. (2019). For its part, TFP growth has slowed in this century. Whether this reflects mis-measurement, a temporary decline, or a new normal (see, e.g., Crafts (2018)) remains to be seen.

²³Authors' calculations based on NIPA data and the Federal Reserve's Financial Accounts.

RP \ RFR	0.0%	-1.0%	-2.0%	-3.0%
2.0 %	-1.8%	-0.6%	+1.7%	+6.0%
3.0 %	-2.0%	-1.3%	+0.3%	+3.2%
4.0 %	-2.1%	-1.7%	-0.7%	+1.4%

Table 1: The impact of defined-benefit pay-go policy on long-run ex-ante utility for different calibration targets for the risk-free rate and the risk-premium.

policies of this size. Since both the current young and current old clearly gain from an introduction of these transfers, increases in EAU for generations born in the long run, which we call *long-run EAU*, indicates, as one would expect and we confirm, a Pareto improvement. Thus, if long-run EAU increases, EAU for all generations rises. Obviously, if long-run EAU falls, the policy is not Pareto efficient.

Table 1 reports, for different risk-free-rate (RFR) and risk-premium (RP) calibrations, the percentage impact on long-run EAU. Clearly, welfare gains increase as the RFR is lowered keeping the RP fixed or as the RP is lowered keeping the RFR fixed. To interpret these results, recall that the model abstracts from both population growth and TFP growth and that rates of return, therefore, have to be regarded as differences relative to the overall growth rate of the economy. If, for instance, these rates add up to 2%, B1 entails a real interest rate of +1% and B2 of 0%. Thus, table 1 shows that higher growth – when keeping the real-world return targets fixed and thereby reducing those targets in the detrended model – provides more leeway for pay-go to Pareto improve.

We next examine our two baselines, B1 and B2, in more detail. Both feature a relatively low RP of 3 percent. B1 calibrates preferences to a -1 percent RFR, while B2 to a -2 percent RFR. For B1, expected utility of those born in the long-run falls by 1.3 percent. For B2, it rises by 0.3 percent. Among the five cases from table 1 which exhibit gains, only B2 is, in our view, remotely plausible.

Policy scale plays an important role in determining long-run EAU impacts. As the left-hand-side (LHS) of figure 1 shows, the percentage change in long-run EAU is negative under B1. But for B2, it starts positive and goes negative at a policy scale equal to roughly 25 percent of the long-run, no-policy, average capital stock. That’s not much larger than the 20 percent value considered by Blanchard (2019). But this policy, which transfers, period by period, 20 percent of average steady-state capital from the young to the old, corresponds to only a 3 (or 4) percent tax on the young’s income in B1 (or B2).²⁴ That’s far below the combined explicit and implicit average wage-tax rate used to finance U.S. intergenerational redistribution. Figure 1 shows major expected utility losses under both B1 and B2 as the economy moves from running pay-go based on a fixed transfer ranging from zero to 25 percent of the young’s no-policy, long-run average of their wages plus endowment.

Figure 2 shows EAU effects of the transfer scheme on both current and prospective generations. Clearly, the initial old gain, as they simply receive a transfer with no strings attached.

²⁴Recall that period length is 25 years. Hence, the yearly transfer is less than 1 percent of aggregate capital.




Figure 1: Long-run EAU impact of pay-go policy, plotted at two different scales, as a fraction of aggregate capital (LHS) and of young's total income (RHS).




Figure 2: Generation-specific EAU impact of pay-go policy along the transition path; initial conditions equal long-run averages.

Their EAU, which we omit from figure 2, increases by 1.8 percent and 2.9 percent in B1 and B2, respectively. The current young also gain substantially, because crowding out takes effect only after they are old. Hence, their wages when young are unchanged, but the rate of return they earn when old on their savings are higher due to the smaller amount of capital they, as a

generation, bring into old age. All later generations gain substantially less (in case of B2) or make outright EAU losses (in the case of B1). Figure 2 also illustrates that the EAU impact on generations born one or two periods after the policy change is already close to the long-run EAU impact. This observation justifies our focus throughout the paper, on long-run EAU.

5 The Crucial Role of Risk-Sharing

This section demonstrates in four ways the crucial importance of risk-sharing to long-run ex-ante utility and, thus, to the prospects for an EAU Pareto improvement. First, we decompose the EAU impact of pay-go in the two baseline calibrations from above, showing that risk-sharing is the source of long-run EAU gains when they arise. It is also a mitigating factor when long-run EAU falls. Second, we show that a policy of two-way transfers, entailing, on average, no intergenerational redistribution, produces larger EAU gains than does pay-go. Third, adding pay-go once this two-way transfer policy is in place does not further Pareto improve. Fourth, we show that Blanchard’s assumed endowment, with its risk-sharing capacity, is key to a Pareto improvement when it arises.

5.1 Decomposing Ex-Ante Utility Gains

To clarify how deficit policy works, we now decompose changes in EAU into non risk-sharing and risk-sharing effects.²⁵ We begin with the *risk-neutral effect* (RNE) referenced in the introduction. It captures the change in EAU that would arise for risk-neutral agents with an IES of 1. Their utility function is defined as follows.

$$\bar{U}_0^t = E_0 \left[c_{y,t}^{1-\beta} \cdot \mathbb{E}_t \{ c_{o,t+1} \}^\beta \right]. \quad (10)$$

The RNE, the ratio between the agent’s utility before and after the introduction of a transfer, picks up the pay-go policy’s crowding out of capital, which, in all of our calibrations, leads, on average, to lower long-run levels of consumption both when young and old.²⁶ In addition, the risk-neutral effect captures any changes in EAU arising from policy-induced changes in the age-consumption profile.

We argue that risk-sharing comes in two distinct forms, one relating to the riskiness of old-age consumption given the date-event of birth, the other to the state into which generations

²⁵For his part, Blanchard (2019) decomposes welfare changes from pay-go as arising from 1) providing agents with a higher safe return than is paid by the safe asset and 2) the crowding out of capital. Blanchard’s equation 3 captures this first effect. His discussion suggests this effect is positive if the safe rate is less than 1. That’s true for the first generation making the transfer. But one needs to average this term over future states of the economy to understand its contribution to the EAU of future generations. Doing so indicates that the expected value of this term equals the sum of a) the product of the average value of X (the difference between 1 and the risk-free rate) and Y (the average value of the marginal utility of second-period consumption) and b) the covariance of X and Y . Both terms depend on risk-sharing arrangements. Hence, Blanchard’s decomposition confounds the impact of pay-go policy on risk-sharing with changes in average consumption values.

²⁶Barbie et al. (2007) provide conditions on prices that ensure that a reduction in investment increases aggregate consumption – conditions that don’t hold in our model.

are born. Recall, the former is called the *life-cycle-risk effect* (LRE) and the latter the *the cohort-risk effect* (CRE). We define LRE as a change in \hat{U}_0^t/\bar{U}_0^t , where the expected ex-interim utility of a generation is given by

$$\hat{U}_0^t = \mathbb{E}_0 \left[c_{y,t}^{(1-\beta)} \mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}^{\frac{\beta}{1-\gamma}} \right]. \quad (11)$$

Yet (11) doesn't capture uncertainty over the state into which one is born. This brings us to CRE, which is defined as the change in U_0^t/\hat{U}_0^t . To sum up, we write ex-ante utility (EAU) as the product of three terms,

$$U_0^t = \bar{U}_0^t \cdot \frac{\hat{U}_0^t}{\bar{U}_0^t} \cdot \frac{U_0^t}{\hat{U}_0^t}, \quad (12)$$

so that percentage changes in EAU equal, to a first order, the sum of percentage changes in the three terms – RNE, LRE, and CRE.

$$\frac{\Delta U_0^t}{U_0^t} \approx RNE + LRE + CRE. \quad (13)$$

Figure 3 decomposes the above reported pay-go policy's long-run EAU changes for B1 and B2 into their RNE, LRE, and CRE components. The RNE effect is, due to the model's crowding out, negative – increasingly so with policy scale. The two risk-sharing effects, LRE and CRE, are, on the other hand, both positive, the LRE effect being greater. RNE, LRE, and CRE sum to the overall impact – the solid blue curve. The reason that curve is not lower, in B1, and positive, for a range, in B2, is thus clearly due to risk-sharing. In both cases the risk-sharing effects are concave, and thus the overall effect is concave as well, which is why in B2 the welfare impact exhibits an interior maximum, around 10 percent, and eventually turns negative, around 25 percent. Thus, our decomposition exercise shows that any welfare gains from pay-go policy reflect improved intergenerational risk-sharing, rather than intergenerational redistribution per se. Note that this conclusion would still stand if we used Blanchard's welfare measure (9) — by definition, the RNE and the LRE are the same for both measures, while the CRE is smaller for his measure.

5.2 Risk-Sharing With Two-Way Transfers

Given the crucial role of risk allocation, we now construct a risk-sharing scheme that does a far better job at allocating risk. Our revised policy transfers from the young to the old if there are two below-median TFP shocks in a row. The old alive after two such shocks are hit with a lifetime double whammy – a particularly low wage when young and a particularly low return to capital when old. In opposite cases, when there are two above-median TFP shocks in a row, the revised policy transfers from the old to the young. In all other states, there is no transfer. Note that this implies that the average net transfer across generations is zero. We refer to this

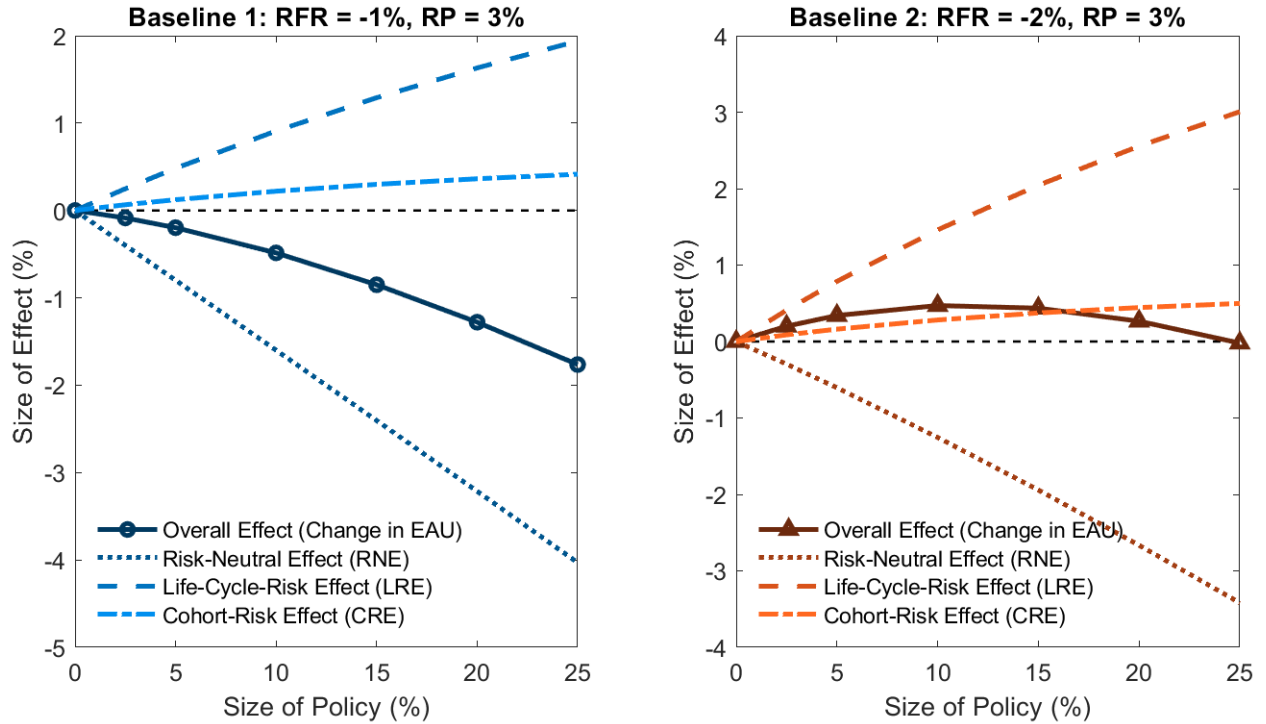


Figure 3: Decomposing the impact of pay-go policy on long-run ex-ante utility.

as the *Two-Way Transfer* (TT) scheme.²⁷

Tables 2 compares and decomposes long-run EAU impacts of pay-go (PG), TT, and a combination of the two (TT+PG): the TT scheme plus a pay-go plan whose transfer is fixed at 5 percent of the no-policy, long-run average capital stock. TT produces a Pareto improvement for both B1 and B2. Indeed, the long-run EAU impacts of TT policy are substantially larger than under pay-go. Moreover, adding even a small pay-go policy on top of the TT policy reduces the values of both B1 and B2 long-run EAUs. Implementing pay-go in the context of TT policy makes early generations better off and future generations worse off. This is true even for cases in which TT leaves the average safe rate negative.

Figure 4 compares DB, TT, and TT+PG and shows their impacts on EAU of current and future generations along the transition.²⁸ Clearly, a transfer scheme designed to share risk is far more efficient than Blanchard's (defined-benefit) pay-go deficit finance, which requires just the right parameters and just the right scale to share risk.

Figure 5 varies the size of the risk-sharing scheme and decomposes the resulting long-run EAU impacts into our three effects. Compared to the pay-go scheme, the negative risk-neutral effect is much smaller in size. This is as expected given that two-way transfers entail no systematic redistribution from the young to the old and, therefore, no systematic crowding out.

²⁷Note that the TT scheme requires transfers from the old generation to the young in case the old experience two consecutive good shocks, earning high wages when young and receiving high capital returns when old, a situation largely descriptive of the baby-boomer generation.

²⁸We assume, as with the pay-go scheme, that the policy is introduced in a period with a median realization of the TFP shock, which implies that there is no immediate transfer. In the subsequent period a below-median or above-median shock triggers half the usual transfer. The current old are thus unaffected. But the young gain because they are, in effect, given an asset for free that hedges against their own old-age consumption risk.

Long-Run EAU Impact (in %)					
Case	RFR	RP	PG	TT	TT+PG
B1	-1.0%	3.0%	-1.3	0.8	0.3
B2	-2.0%	3.0%	0.3	2.3	2.1

Decomposition of Long-Run EAU Impact (in %)									
Case	RNE			LRE			CRE		
	PG	TT	TT+PG	PG	TT	TT+PG	PG	TT	TT+PG
B1	-3.2	-0.9	-1.6	1.6	1.7	2.0	0.4	0.0	-0.1
B2	-2.7	-0.8	-1.4	2.6	2.7	3.3	0.4	0.4	0.3

Table 2: Comparing and decomposing long-run EAU impacts from pay-go (PG), two-way transfers (TT), and two-way transfers plus 5% pay-go (TT+PG) policies.

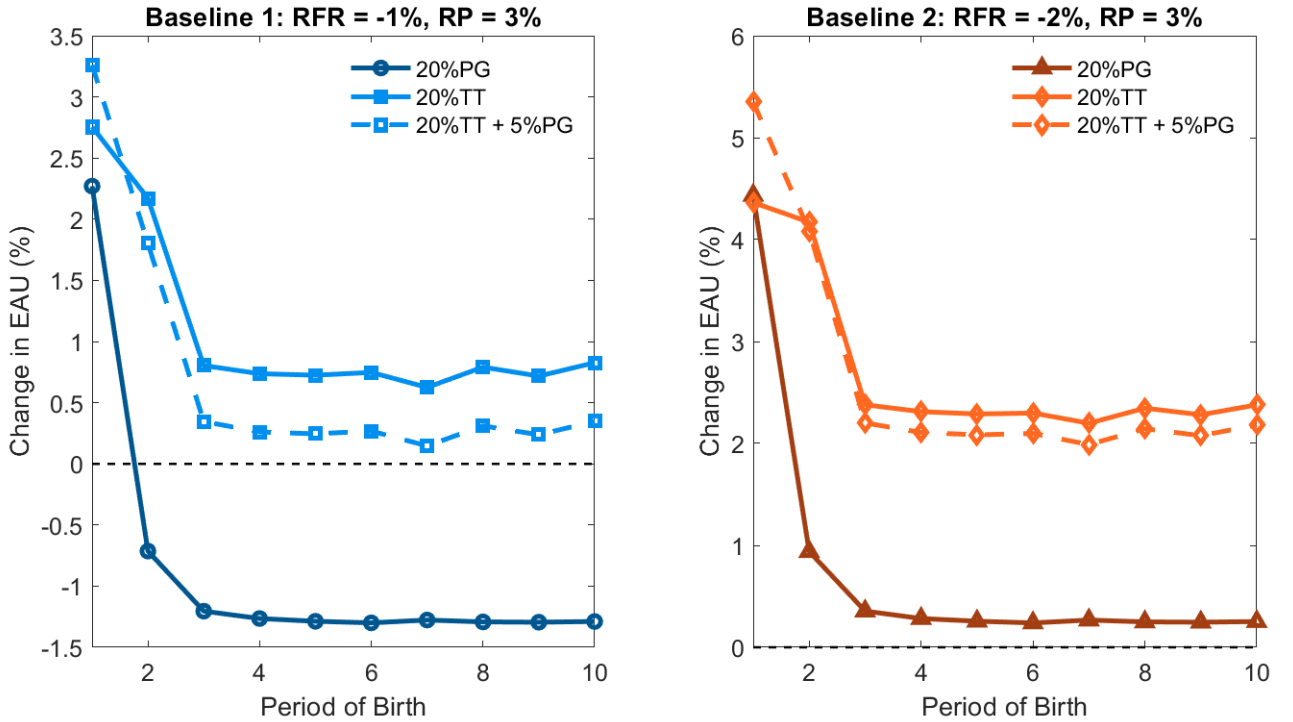


Figure 4: Generation-specific EAU impact of different policies along the transition path; initial conditions equal long-run averages.

As for the life-cycle-risk effect, it slightly improves relative to the pay-go case. The cohort-risk effect is smaller or similar depending on the calibration. Thus, this policy achieves a similar level of risk-sharing between cohorts with much less crowding out, resulting in much larger increases in EAU overall. Comparing figures 3 and 5 reveals that under the B1 calibration, pay-go never Pareto improves, whereas TT always does. This constitutes further evidence that risk-sharing, not systematic redistribution from the young to the old is of central and, potentially, exclusive importance.

Our two-way transfers scheme is certainly sub optimal. As section 3.3 suggests, the optimal scheme would surely depend non-linearly on the economy's state vector – its TFP and stock of capital. But the above analysis demonstrates the ample room for Pareto-improving policy that

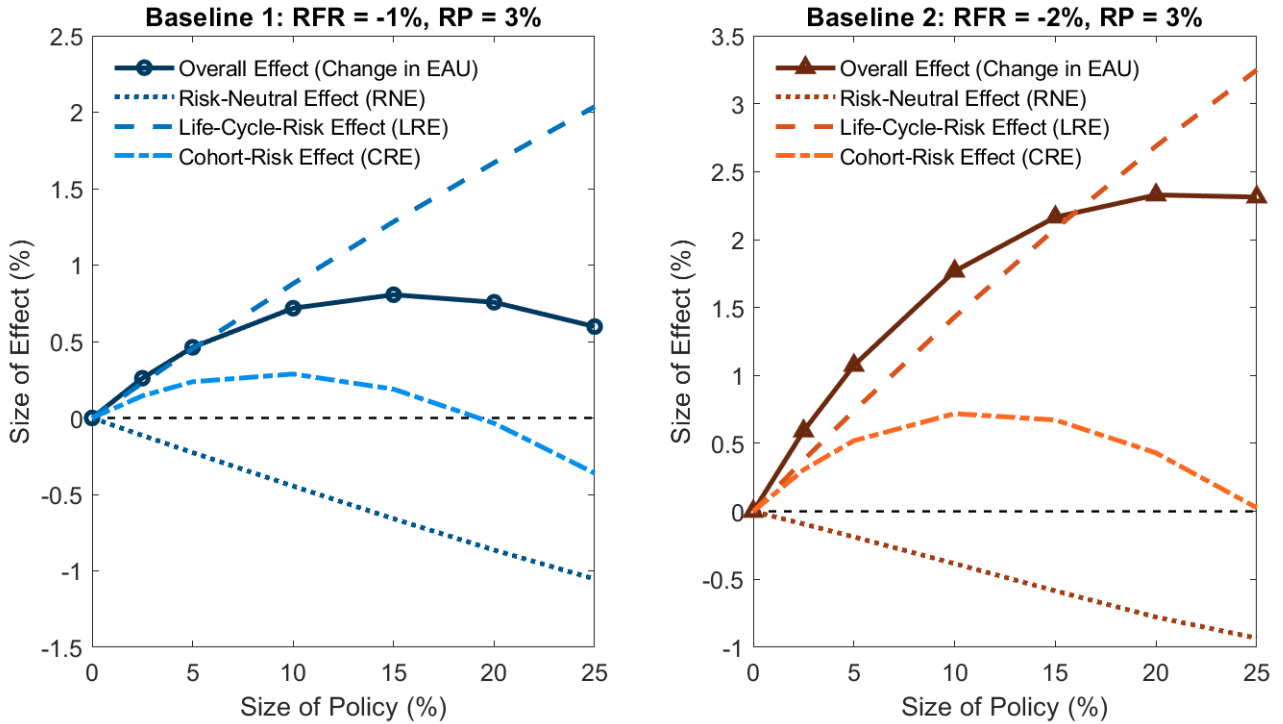


Figure 5: Decomposing the impact of two-way transfer (TT) policy on long-run EAU.

doesn't systematically redistribute. Moreover, as figure 4 shows, when even our crude TT risk-sharing scheme is in place, adding pay-go policy scaled at 5 percent reduces the EAU of future generations – even though the average risk-free rate stays below zero when that additional policy is introduced.²⁹ Indeed, only the initial old benefit from the addition of Blanchard's policy. This result, again, indicates that deficits are desirable only insofar as they help share risk across generations. If risk is already well shared, intergenerational redistribution, whether run under the heading deficit finance, structural tax change, pay-go social security, or something else, will benefit early generations at a cost to future generations.

Finally, note that under TT, the risk-free rate averages, on an annualized basis, -0.4 percent (-1.3 percent) in B1 (B2). In comparison, the mean of the RFR is 0.6 percent (-0.2 percent) in periods when transfers from the young to the old are positive and -1.3 percent (-2.3 percent) when they are negative. This substantial difference reflects three factors. First, low saving and investment in the prior period when the TFP shock was bad. Second, low saving and investment in the current period when the TFP shock is bad. These effects are both present even without policy, yet there is now a third mechanism that reinforces the resulting counter-cyclicality of interest rates – the reduced demand for the safe asset by the young who know they will receive a transfer when old if another bad TFP shock is to come. Thus, loose fiscal policy should coincide with high RFRs and tight fiscal policy should coincide with low RFRs. Stated differently, a high, not a low RFR is the time to run deficits and a low, not a high RFR is the time to run surpluses.

²⁹In B1, the RFR increases from -0.4 percent to -0.1 percent, whereas in B2 it goes from -1.3 percent to -0.9 percent.

RP \ RFR	0.0%	-1.0%	-2.0%	-3.0%
2.0 %	-3.8% (-1.8%)	-2.9% (-0.6%)	-0.8% (+1.7%)	+3.5% (+6.0%)
3.0 %	-4.3% (-2.0%)	-4.0% (-1.3%)	-2.8% (+0.3%)	-0.1% (+3.2%)
4.0 %	-4.4% (-2.1%)	-4.4% (-1.7%)	-3.8% (-0.7%)	-2.3% (+1.4%)

Table 3: Impact of pay-go policy on long-run EAU in a model without the fixed endowment of the young (results from the model with the endowment are in brackets).

5.3 Role of the Endowment

While wages and capital income, in the context of 100 percent depreciation of capital in each period, are perfectly correlated via the TFP shock, the fixed endowment makes the resources of the young (the endowment plus their wages) less risky than that of the old (their capital income). Unfortunately, the young can't, in their infancy, make risk-sharing deals with their parent's generation. This creates a missing market, which the government can implicitly emulate via policy.

To quantify the endowment's importance, we simulate pay-go in the closed economy but without the endowment.³⁰ Table 3 shows the long-run EAU impact from implementing pay-go in the no-endowment economy for different calibration targets. Cell-specific results for the economy with the endowment are in brackets. The long-run EAU impact is positive in only one case – with a negative risky return.³¹ Otherwise, there are welfare losses that are often significant. Compared to the model with the endowment, the long-run EAU impact is roughly two percentage points worse for most calibrations.

To illustrate the very different message that the economy without the fixed endowment sends, figure 6 plots the long-run EAU impact as a function of the calibration target for the risk-free rate. It does so in three ways. We first reconsider the economy with the fixed endowment and lower risk-free rates (going from right to left), while keeping the risk premium calibration target fixed. This curve, despite being mostly in negative territory, indicates that transfers are more desirable the lower the risk-free rate. This message changes however, when, instead of the risk premium, the risky rate is held fixed (which is achieved by increasing the risk-premium via an increase in the risk-aversion parameter). In this case the curve becomes substantially flatter and does not reach positive territory even for a risk-free rate of minus three percent. Finally, we move to the economy without a fixed endowment. Now reducing the risk-free rate no longer improves long-run EAU whatsoever. In short, a low risk-free rate does not, per se, represent a general invitation to run deficits.³²

³⁰In so doing, we truncate the TFP-shock to avoid potential transfer-scheme collapse. Truncating at nine standard deviations suffices for this purpose. Second, we adjust capital's share in the production function to ensure that the new model's capital share matches Blanchard's effective capital share. This adjustment doesn't materially affect our results.

³¹While table 3 shows that welfare gains are hard to obtain when the risky return is positive, it is not impossible: with a RFR of -3.5% and a RP of 3.6% (i.e. risky return of 0.1%) EAU improves by 1.0% .

³²This point aligns with those made in Brumm et al. (2022).




Figure 6: Long-Run EAU impacts of pay-go for different calibration targets for RFR and RP. The top curve changes the RFR target while keeping RP fixed. The middle curve changes the RFR target while keeping the risky return fixed. The bottom curve changes the RFR target while keeping the risky return fixed, but now in a model without the endowment assumption.

6 Open-Economy Findings

As is clear, the reason deficit finance makes future generations worse off, under reasonable calibrations, is its crowding out of capital. But in an open-economy, domestic saving reductions are spread globally. This limits domestic crowding out, leaving deficit finance more potential to Pareto-improve ex-ante utility of domestic residents – current and future – via enhanced risk-sharing. However, domestic Pareto gains come at a price to foreigners as they have less capital with which to work and, thus, earn lower wages. In this case, putative domestic Pareto improvements partly reflect beggar-thy-neighbor policy. This said, foreigners may also benefit from improved international risk-sharing. To carefully assess these issues we add a foreign country to our model.

6.1 Calibration

We now consider a calibration that treats the U.S. as the domestic country and the rest of the world (RoW) as the foreign country. We re-calibrate the model to match the relative sizes of the two economies, the volatility and correlation of their productivity shocks, and the sizes of international portfolio positions. Appendix F provides calibration details. The basic targets are as follows. Based on data from the Penn World Tables, RoW GDP is 6 times U.S. GDP and 1.25 times more risky. The correlation of the TFP shocks of the two countries is 0.22 and their auto-correlation is, as in the closed-economy model, assumed to be zero. We choose identical preferences parameters for the two regions – parameters that match the risk-free rate and the

RP \ RFR	0.0%	-1.0%	-2.0%	-3.0%
	U.S. / RoW	U.S. / RoW	U.S. / RoW	U.S. / RoW
2.0 %	-0.6% / -0.1%	+0.4% / -0.0%	+2.4% / +0.2%	+6.3% / +0.6%
3.0 %	-0.7% / -0.1%	+0.0% / -0.1%	+1.5% / +0.1%	+4.3% / +0.4%
4.0 %	-0.7% / -0.2%	-0.2% / -0.1%	+0.8% / -0.0%	+2.9% / +0.2%

Table 4: Long-run EAU impact of domestic pay-go policy for different calibration targets.

equity premium in the home country, i.e. the U.S. The 2019 data from the Bureau of Economic Analysis reports RoW-held U.S. capital worth 80 percent of U.S. GDP, U.S.-held RoW capital worth 85 percent of U.S. GDP, and net U.S. bond holdings worth 40 percent of U.S. GDP.³³ To fit these data, we introduce country-specific compensated investment costs that reduce the return to cross-country investing. The terms δ_H , δ_F , and δ_{FB} reference the cost to domestic agents of investing in foreign capital, the cost to foreigners of investing in home capital, and the cost to foreigners of investing in the bond. Our calibration uses these three parameters to match the three cross-country asset positions, $(k_{H,F}, k_{F,H}, b_H)$.³⁴ In B1 the annualized cost parameters needed to match the targets are: a 0.9 percent cost of the U.S. investing in RoW capital, a 2.9 percent cost of RoW investing in U.S. capital, and a 1.1 percent cost to RoW of investing in the international bond. With these costs, the associated risk-aversion coefficient needed to match the risk premium is 21.4.³⁵ Additional details of the calibration can be found in appendix F.

6.2 Pay-Go — More Favorable Due to Openness

Given this calibration, we now consider the long-run EAU consequences of a pay-go policy introduced in the home country, which turns out to be only slightly more favorable than in the closed economy case. As table 4 shows, realistic values of the risk-free rate and the risk-premium produce long-run EAU losses. The introduction of a transfer scheme in the U.S. has negative long-run EAU effects for the rest of the world for realistic values of the return to capital. At the same time this slightly alleviates the negative effect in the home country. There are now several cases (e.g. a risk-free rate of -1 percent with a risk premium of either 2 or 3 percent) where the home country enjoys (modest) long-run EAU gains whereas the foreign country experiences losses. In this case, the beggar-thy-neighbor effect is strong enough to imply long-run EAU gains in the home country that would not materialize in a closed economy. Finally, enhanced international risk-sharing can actually lead to situations where both the home and foreign country gain, albeit for highly unrealistic values of the average return to capital.

³³Capital in the model corresponds to equity and foreign direct investment from the data, and bond holdings in the model correspond to net debt securities – debt securities held minus debt liabilities.

³⁴Given our assumed identical preferences in both regions, including unitary intertemporal elasticities of substitution, we have two preference parameters and three cost parameters to match the risk-free rate, risk premium in the home country, and the asset positions of the two countries in the context of global bond-market clearing, i.e., we have five parameters to match five moments.

³⁵Under B2, the calibrated cost parameters are slightly higher and the risk-aversion coefficient is slightly lower.

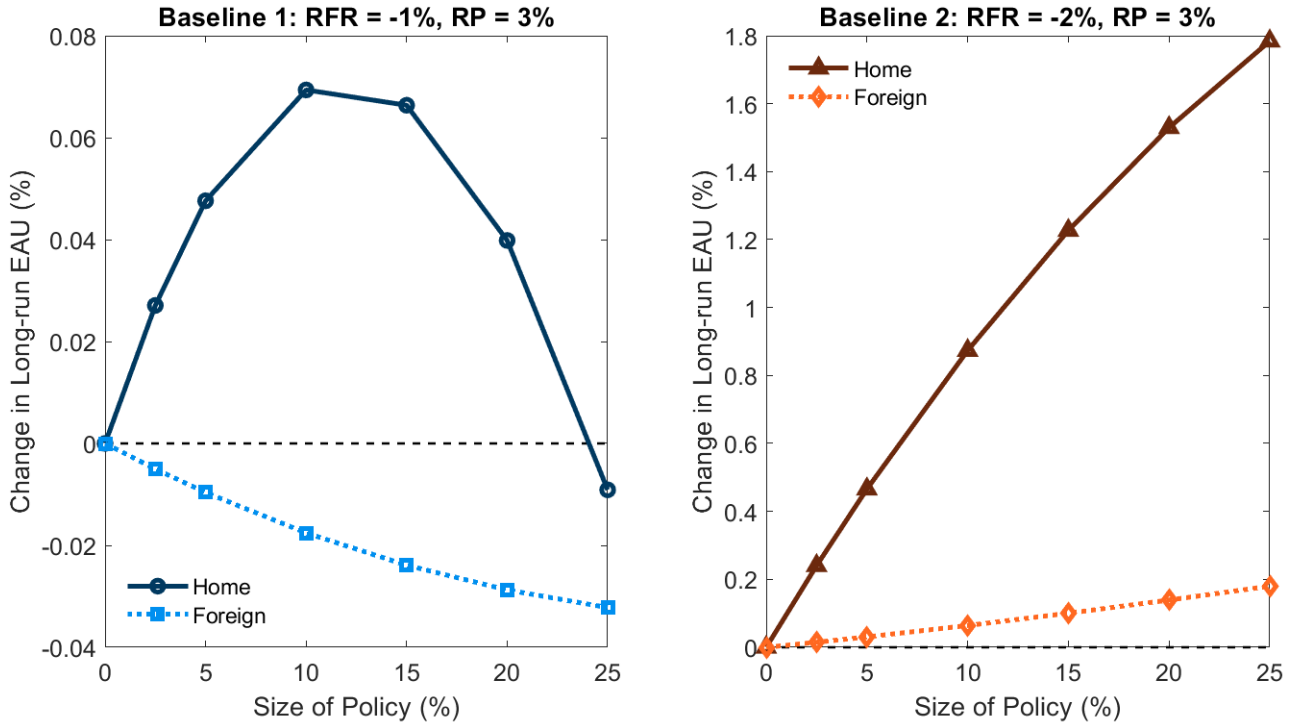


Figure 7: Long-run EAU impacts of domestic pay-go policy in the two-country model under calibrations B1 and B2.

To illustrate this in more detail, we now consider the two baseline calibrations from our closed-economy analysis. The left-hand panel of figure 7 confirms the intuition. While in a closed economy, long-run EAU gains were impossible under B1, the home country now experiences small gains, and yet the rest of the world experiences losses. The right-hand panel of figure 7 illustrates that in B2, domestic long-run EAU improves significantly for a large range of transfer payments, while the rest of the world now also gains.

The importance of international risk-sharing becomes clearer when we decompose long-run EAU effects. The upper two panels in figure 8 decompose the long-run EAU effects in the home and foreign countries for B1. While the long-run EAU effect for the foreign country is negative due to crowding out, both risk-sharing effects are actually significantly positive. This is confirmed in the lower panel of the figure that considers B2. Here, both countries gain, with the welfare improvement clearly driven by risk-sharing.³⁶ Since the home government is not redistributing among foreign generations, the only source of the improved risk-sharing is international risk-sharing through the bond market. The reduction in domestic saving reduces the domestic demand for safe bonds, lowering their price and raising their return. This permits foreign agents to achieve a safer portfolio at a lower cost. Still, achieving a global Pareto improvement requires invoking, in the case of B2, a rather low risk-free rate and an unrealistically low value for the risk-premium.

To sum up, depending on calibration targets we observe three cases. For realistic returns

³⁶Note, however, that the beggar-thy-neighbor effect now even implies a positive RNE for the home country. While the crowding out of domestic capital still has a negative effect, that effect is now smaller than the positive effect of a shift in resources from young to old. For context, see the discussion of the RNE in section 5.1.

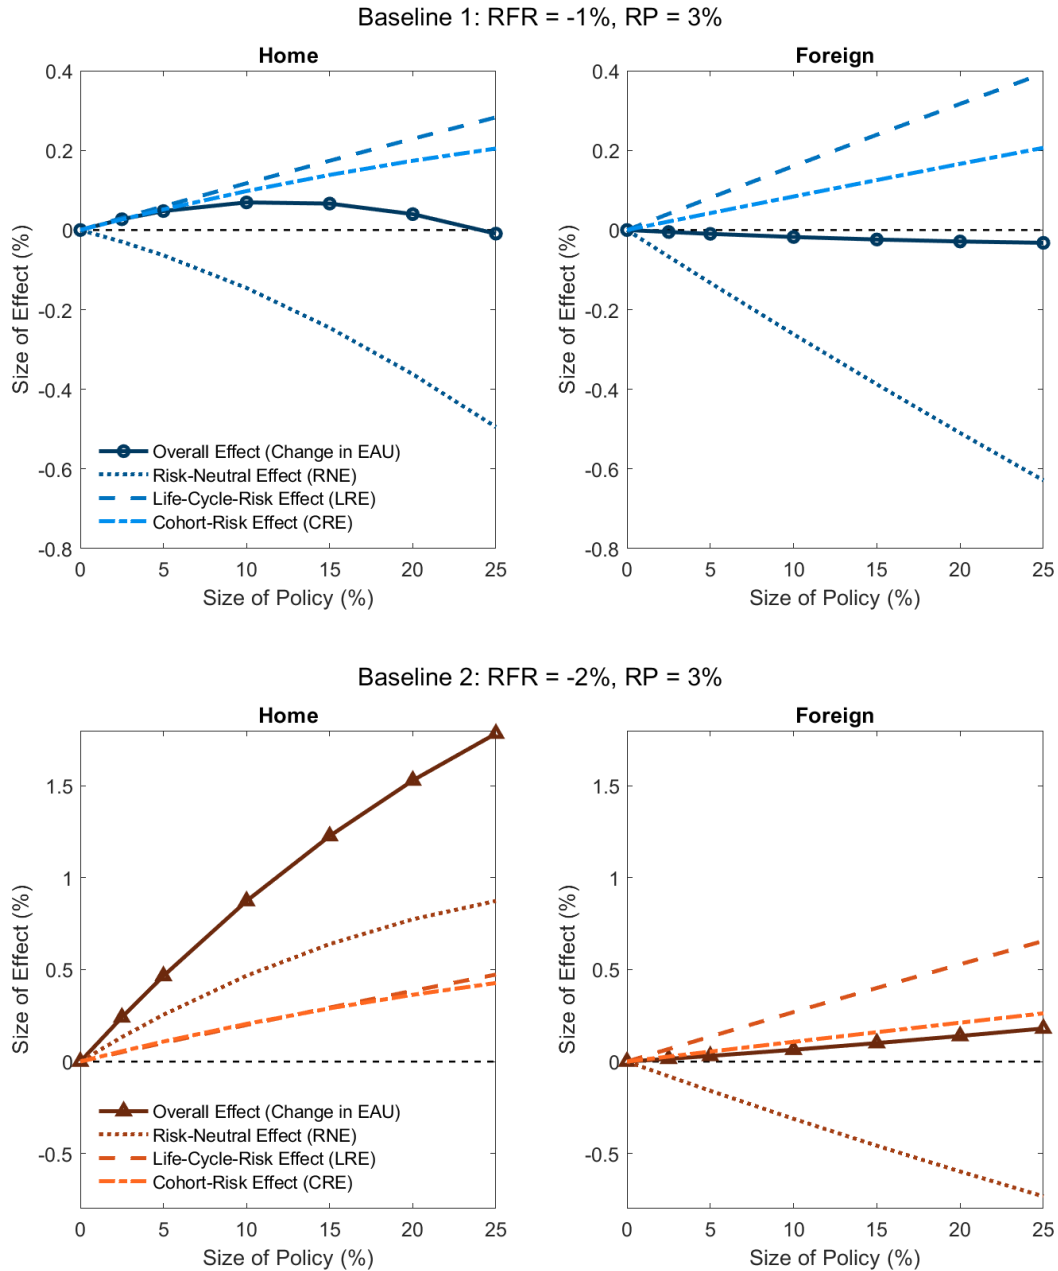


Figure 8: Decomposing long-run EAU impacts of domestic pay-go policy in the two-country model for calibrations B1 and B2.

to capital, domestic and foreign long-run EAU both decline. For very low returns to capital, both countries gain from deficit finance in the home country. For cases in the middle, the beggar-thy-neighbor effect implies home-country gains at the price of foreign-country losses.

6.3 Two-Way Transfers — Superior Once Again

So far, we have learned that pay-go policies have a somewhat better chance to Pareto-improve domestic agents when an open economy is considered. Yet we have also seen that such potential gains still stem from risk-sharing, while the counteracting risk-neutral effect is dampened due to beggaring-thy-neighbor. These findings suggest that implementing the two-way transfer scheme

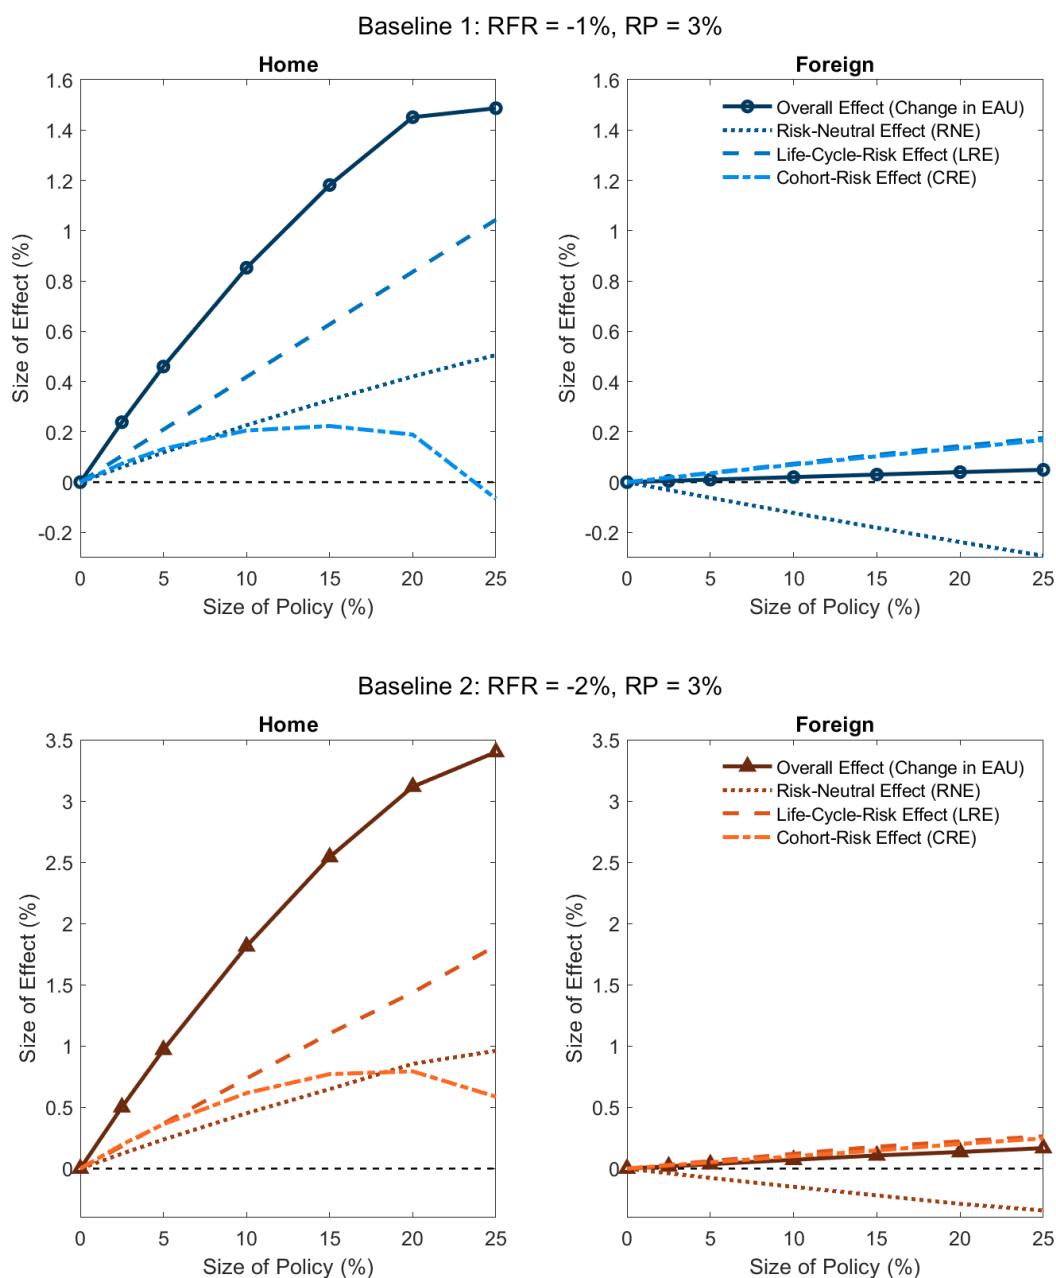


Figure 9: Decomposing long-run EAU impacts of domestic two-way transfer policy in the two-country model for calibrations B1 and B2.

in the open economy might be superior to the pay-go policy once again, and that it might have even larger benefits than that same scheme when implemented in a closed economy. Indeed, figure 9 confirms both conjectures. It decomposes the long-run EAU impacts of two-way-transfer policy implemented in the home country on the home and foreign country for varying policy sizes. Compared to pay-go, such a state contingent policy is much more beneficial to home country residents, as the risk-neutral and life-cycle-risk effects are more favorable. Foreigners, in turn, enjoy comparable risk-sharing benefits at much lower risk-neutral losses and thus experience higher EAU gains, which are now positive even under B1. Thus, the two-way-transfer policy in the open economy underscores the point that long-run EAU gains in this model are about enhanced risk-sharing, both between generations and countries.

7 Limitations of the Model

We examine a major policy question in a very simple and stylized model taken from Blanchard (2019) in both closed- and open-economy setups. First, our analysis considers a single representative agent per generation and country. However, individuals that differ in preferences and endowments will be affected differently by a social security system. For example, we equate the safe rate with returns on short-term Treasuries. Yet as pointed out by Brumm et al. (2020), close to 90 percent of Americans are in debt and their safe real rates – the safe real rates they can earn by pre-paying their mortgages, credit-card balances, etc. – are much higher and thus equal or exceed the real growth rate. Taking this into account makes social security a much less attractive proposition for a large part of the U.S. population. Furthermore, the issue of under-accumulation of capital can be substantially more detrimental to welfare in the presence of idiosyncratic shocks – see Davila et al. (2012). On the other hand, as pointed out by Imrohoroglu et al. (1995), a pay-go social security system can help share intragenerational risk, which is obviously absent from our analysis.

Second, our focus, like Blanchard’s, is purely on government redistribution among generations. We are not considering deficit finance used to fund infrastructure, provide public goods, or correct externalities, like global warming. Nor do we consider the value of deficits as counter-cyclical policy. Our model also abstracts from monetary aspects of the economy. It therefore omits an important problem that low real interest rates can cause, namely making it more likely that nominal interest rates hit their effective lower bound. Under such circumstances, monetary policy might find it hard to stimulate the economy and output may fall below the full employment level. Fiscal policy, in contrast, might be able to sustain demand and in addition alleviate the problem of the effective lower bound by raising real rates (see, e.g., Eggertsson et al. (2019)).³⁷

Third, existing pay-go policy matters. As we’ve shown, the impact on long-run EAU either steadily declines or rises and then declines with policy scale. Hence, calibrating and running Blanchard’s model assuming, as he does, that the U.S. has no initial policy in place appears to give pay-go an unwarranted efficiency advantage.

Finally, this paper follows Blanchard (2019) in adopting a quite narrow view on the reasons interest rates are low, thereby abstracting from some reasons for low rates that might favor higher levels of public debt. These include idiosyncratic risk as in Aiyagari and McGrattan (1998), Brumm et al. (2022), and Reis (2021), income inequality in the presence of non-homothetic preferences as in Mian et al. (2021b), or convenience benefits of government debt as in Mian et al. (2021a).

³⁷Moreover, in a multi-country setup where several countries struggle with the effective lower bound, raising debt or social security in one country can have – in addition to the negative beggar-thy-neighbor spillover and the positive risk-sharing channel we identify – a positive spillover by elevating the problem of the lower bound in the other countries, as modeled in Eggertsson et al. (2016).

8 Conclusion

The decades-old deterministic literature on the potential over-accumulation of capital raises a green flag when the marginal product of capital falls below the growth rate. The flag signals the ability to perpetually take from the young and give to the old, making all generations better off. Although measurements differ, there is, unfortunately, no doubt that the marginal product of capital exceeds the growth rate most of the time.³⁸

With uncertainty, the return to safe assets – the safe rate – allegedly comes into play. The U.S. safe rate routinely runs below the U.S. growth rate. This has led many economists to treat the average difference between the two as an arbitrage opportunity – one that can also be exploited via an ongoing pay-go policy. Of course, economics teaches us how to decide whether a given policy would, on balance, benefit some or all generations without harming others. Blanchard (2019)’s thought-provoking paper suggests that conditions may be right, in the U.S. and other low-interest-rate countries, to run fiscal deficits at no cost. No cost means helping current old generations at no cost to their progeny. We examine the “deficits are free” proposition using Blanchard’s model, but with an important correction to his formula for ex-ante utility, which leaves even more room to find Pareto improvements from deficit finance.

Blanchard’s framework does, indeed, admit cases for which deficit finance can Pareto improve from an ex-ante perspective. But the calibrations required for such an outcome stretch reality. Under close inspection, it is clear why low safe rates are not a clear invitation to run deficits. Low safe rates signal a strong demand for safety. This need is directly addressed via systematic risk-sharing, not intergenerational redistribution. Blanchard (2019)’s model is a two-period, OLG model with aggregate risk. In such models, Pareto-improving, government-organized, intergenerational risk-sharing policy is ripe for the picking given the inability of the living to trade with the unborn. Our decomposition of policy impacts on ex-ante utility into risk-neutral and risk-sharing factors identifies the potential Pareto gains available from risk-sharing. In contrast, even crude, government-organized intergenerational risk-sharing – transfers running from the young to the old and vice versa depending on the economy’s current and prior states – can materially improve current and future generations’ ex-ante utility. As for deficits, when they Pareto improve, they do so serendipitously – not because they are designed to solve the problem at hand, namely addressing missing private financial markets, but because they too can share risk under very special conditions. But, as we show, once proper risk-sharing policy is in place, deficit finance provides no further scope for sharing risk.

We also explore deficit finance in an international context. In such settings, the crowding out of domestic saving translates less than one for one into reduced domestic investment. This limits the reduction in domestic wages, albeit at the price of lower foreign wages. On the other hand, moving to an open economy raises the prospect of mutually beneficial international risk sharing. Yet, here again, improved risk-sharing, in this case, both between generations and

³⁸Accordingly, in all our calibrations above, the unconditional mean of the risky rate of return, R , which equals the marginal product of capital in the model, substantially exceeds the average growth rate, which is normalized to zero. This also holds true for the geometric mean, which Zilcha (1991) points out is the relevant statistic.

between countries, not systematic redistribution from the young to the old or from foreign to domestic residents, may be best suited for achieving a Pareto improvement. In sum, our paper's message is clear. When interest rates go low, government risk-sharing, not deficit finance, provides the likely path to economic efficiency.

APPENDIX

A Characterizing Efficient Allocations

We consider allocations that are solutions to the following planning problem

$$\max_{(c_{o,t}, c_{y,t}, S_t)_{t=0}^{\infty}} \sum_{t=0}^{\infty} \xi_t E_0 U(c_{y,t}, c_{o,t+1}) \quad (14)$$

$$\text{s.t. } c_{y,t} + c_{o,t} = f(A_t, S_{t-1}) - S_t, \quad \forall t$$

where $S_t = K_{t+1}$ denotes aggregate savings and multipliers are assumed to be summable, $(\xi_t)_{t=0}^{\infty} \in \ell_1$. As Dechert (1982) shows, this problem can arise from a constrained optimization problem where the utility of one generation is maximized keeping all other generations at specified utility levels, including status-quo utility levels. Solutions to problem (14) are clearly Pareto efficient.

We derive necessary conditions on allocations that solve (14) for the model in section 3.1, which features additive separable utility, $U(c_y, c_o) = (1-\beta)v(c_y) + \beta\mathbb{E}[v(c_o)]$, and Cobb-Douglas production. The first order condition with respect to $c_{y,t}$ and $c_{o,t}$ implies

$$\frac{\xi_t}{\xi_{t-1}} \frac{(1-\beta)v'(c_{y,t})}{\beta v'(c_{o,t})} = 1.$$

Efficient allocations are characterized by a sharing rule that is time-varying, but not state-varying. If agents' utility exhibits constant relative risk aversion, $v(c) = \frac{c^{1-\gamma}}{1-\gamma}$, the sharing rule is linear – see Proposition 1(a). Hence, in any efficient allocation, at a given t , the young and the old consume a fixed fraction of aggregate consumption, independent of the history of shocks. In the following, we denote the fraction of the young's consumption at t by λ_t . Direct computation gives

$$\lambda_t = \frac{((1-\beta)\xi_t)^{\frac{1}{\gamma}}}{((1-\beta)\xi_t)^{\frac{1}{\gamma}} + (\beta\xi_{t-1})^{\frac{1}{\gamma}}}. \quad (15)$$

For $E > 0$, this cannot be satisfied in a laissez-faire equilibrium – see Proposition 1(b). Moreover, it is always possible to make all future generations strictly better off by guiding them to the appropriate efficient risk-sharing allocation. This is independent of the risk-free rate.

We now characterize the planner's optimal saving function, $S_t(A, k)$, by taking the first order conditions of (14) with respect to savings. We obtain

$$\begin{aligned} & (1-\beta)\lambda_t^{-\gamma}(E + f(A, k) - S_t(A, k))^{-\gamma} \\ &= \beta(1-\lambda_{t+1})^{-\gamma} \sum_{A'} \pi_{A'} R(A', S_t(A, k)) (E + f(A', S_t(A, k)) - S_{t+1}(A', S_t(A, k)))^{-\gamma}. \end{aligned}$$

Note that this is identical to the Euler equation of an infinitely lived representative agent

with the same CRRA cardinal utility but with a time varying instantaneously discount factor $\frac{\beta}{1-\beta} \left(\frac{\lambda_t}{1-\lambda_{t+1}} \right)^\gamma$. If this is assumed to be (uniformly) less than one, an optimal solution always exists, and that completely characterizes the solution to (14).

Proposition 1 and its proof make the above points more explicit. To recapitulate, we denote a history of shocks by $A^t = (A_0, \dots, A_t)$, reference aggregate consumption at any date event A^t by $C(A^t)$, and restate Proposition 1.

Proposition 1: *Suppose agents' utility exhibits constant relative risk aversion.*

(a) *For all t , there is, for any Pareto-efficient allocation, a λ_t , for $\lambda_t \in [0, 1]$, such that*

$$c_y(A^t) = \lambda_t C(A^t), \text{ for all } A^t.$$

(b) *For $E > 0$, no laissez-faire equilibrium is Pareto efficient.*

(c) *For any sequence $(\lambda_t)_{t=0}^\infty$ with $\lambda_t \in (0, 1)$ and $\frac{\beta}{1-\beta} \left(\frac{\lambda_t}{1-\lambda_{t+1}} \right)^\gamma < 1$, for all t , there is a Pareto efficient allocation that can be implemented as a competitive equilibrium with lump-sum transfers by a time-varying transfer scheme*

$$T_t(k, A) = (1 - \lambda_t) (E + W(A, k) - S_t(A, k)) - \lambda_t k R(A, k),$$

with the savings functions $S_t(A, k)$ solving

$$\begin{aligned} & (1 - \beta) \lambda_t^{-\gamma} (E + f(A, k) - S_t(A, k))^{-\gamma} \\ & = \beta (1 - \lambda_{t+1})^{-\gamma} \sum_{A'} \pi_{A'} R(A', S_t(A, k)) (E + f(A', S_t(A, k)) - S_{t+1}(A', S_t(A, k)))^{-\gamma}. \end{aligned}$$

Proof. Fixing any time period, t , given aggregate consumption, $C(A^t)$, at nodes A^t , a necessary condition for Pareto-efficiency is that there are $\xi_{t-1} > 0$ and $\xi_t > 0$ such that

$$\begin{aligned} (c_y(A^t), c_o(A^t))_{A^t} & \in \arg \max \sum_{A^t} \pi(A^t) (\xi_t (1 - \beta) v(c_y(A^t)) + \xi_{t-1} \beta v(c_o(A^t))) \\ \text{s.t. } & c_y(A^t) + c_o(A^t) = C(A^t) \quad \forall A^t \end{aligned}$$

The first order conditions for this optimization problem imply that

$$\frac{\xi_t}{\xi_{t-1}} \frac{(1 - \beta) v'(c_y(A^t))}{\beta v'(c_o(A^t))} = 1 \text{ across all date events } A^t.$$

Using market clearing, $c_y(A^t) + c_o(A^t) = C(A^t)$, we can solve for optimal individual consumption as a function of the multipliers and aggregate consumption. If agents' utility exhibits constant relative risk aversion, we obtain that efficient allocations are characterized by a linear sharing rule, and hence in any efficient allocation, at a given t , the young and the old consume a fixed fraction of aggregate consumption, independently of the history of shocks. In a laissez-faire

equilibrium, by the budget constraints, consumption when young and old can only be colinear if $E = 0$; therefore, for $E > 0$ laissez-faire equilibria are always inefficient.

We assume that the agent takes transfers as given and chooses optimal savings to maximize utility. Optimal transfers are given by

$$T_t(A, k) = (1 - \lambda_t) (E + W(A, k) - S_t(A, k)) - \lambda_t k R(A, k).$$

The young's first order condition for optimal savings reads as

$$\begin{aligned} & (1 - \beta) \lambda_t^{-\gamma} (E + f(A, k) - S_t(A, k))^{-\gamma} \\ &= \beta (1 - \lambda_{t+1})^{-\gamma} \sum_{A'} \pi_{A'} R(A', S_t(A, k)) (E + f(A', S_t(A, k)) - S_{t+1}(A', S_t(A, k)))^{-\gamma}. \end{aligned}$$

It is easy to see that the resulting competitive equilibrium (with transfers) is Pareto-efficient. As in Ball and Mankiw (2007) we can imagine an economy where all agents trade ex-ante in a complete set of Arrow securities. With the prescribed transfers, equilibrium trades are zero. The assumption that $\frac{\beta}{1-\beta} \left(\frac{\lambda_t}{1-\lambda_{t+1}} \right)^\gamma < 1$ for all t implies that prices are summable, and Pareto-efficiency of the allocation follows from a standard argument (see, e.g., Barbie et al. (2007)). \square

B Details on Computation

We first describe our computational approach for the closed economy model and then turn to the two-country case. The state of the closed economy at time t is characterized by capital, k_t , accumulated in the previous period and TFP, A_t , determined exogenously. These variables jointly determine factor prices and consumption of the old,

$$\begin{aligned} R_t &= \alpha A_t k_t^{\alpha-1}. \\ W_t &= (1 - \alpha) A_t k_t^\alpha. \\ c_{o,t} &= k_t R_t + T_t. \end{aligned} \tag{16}$$

Given the state and prices in t , current choices of the young and the risk-free rate, $(c_{y,t}, k_{t+1}, R_{t+1}^f)$, satisfy the following system of equilibrium conditions – two Euler equations of the young, and their budget constraint:

$$\begin{aligned} \frac{1 - \beta}{c_{y,t}} &= \beta \frac{\mathbb{E}_t \{ R_{t+1} c_{o,t+1}^{-\gamma} \}}{\mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}}, \\ \frac{1 - \beta}{c_{y,t}} &= \beta \frac{R_{t+1}^f \mathbb{E}_t \{ c_{o,t+1}^{-\gamma} \}}{\mathbb{E}_t \{ c_{o,t+1}^{1-\gamma} \}}, \\ c_{y,t} &= W_t + E - k_{t+1} - T_t, \end{aligned} \tag{17}$$

where the risky return, R_{t+1} , and consumption when old, $c_{o,t+1}$, depend, as in equation (16), on capital saved for next period, k_{t+1} , and on next period's realization of TFP, A_{t+1} :

$$\log A_{t+1} = \epsilon_{t+1}, \quad \epsilon_{t+1} \sim_{i.i.d.} \mathcal{N}\{0, \sigma^2\}. \quad (18)$$

We solve our closed- and open-economy models on a period-by-period basis for 1000 periods. When the economy is closed, the solution devolves to solving, for each period, three equations in three unknowns – the risk-free rate, the consumption of the young, and the saving of the young. In the open economy, there are seven equations in seven unknowns – the risk-free rate, the consumption and saving of the young in each economy, home-country holdings of foreign capital, and foreign-country holdings of domestic capital. We find exact solutions to the relevant equations using a non-linear solver and determining expected values via Gauss-Hermite quadrature of order 20.

C Details on Ex-Ante Utility

This appendix provides details for understanding our EAU measure: first, comparing it to Blanchard's measure; second, giving a simple example of how it works; third, providing a formal derivation.

C.1 Comparison of EAU Measures

Figure 10 compares long-run EAU effects under both our and Blanchard's EAU measures for our two baseline calibrations in the closed economy. While the effects are qualitatively similar under both measures, it is clear that with our measure, EAU effects are more favorable.

C.2 Simple Example of Ex-Ante Utility

Consider the problem of evaluating, at time 0, the expected utility of an agent born at time 1 who lives for two periods and has Epstein-Zin preferences as specified in (6). Suppose there are two equally likely states at time 1, A and B. Evaluated at time 1, the agent's utility, conditional on state A, is

$$U_1^A = c_1^{(1-\beta)} \mathbb{E}_A [c_2^{1-\gamma}]^{\frac{\beta}{1-\gamma}}, \quad (19)$$

and similarly for state B. Given this measure, how should we evaluate the agent's welfare at period 0 when the agent is not yet born and the state in period 1, A or B, is still uncertain? One option is to simply take the expected value of time-1 utility, namely

$$\hat{U}_0 = 0.5U_1^A + 0.5U_1^B. \quad (20)$$

We call this measure expected ex-interim utility. Now suppose $c_1^A = c_2^A = 1$ (implying $U_1^A = 1$) and $c_1^B = c_2^B = 3$ (implying $U_1^B = 3$) and consider a policy that, if introduced at time zero, will

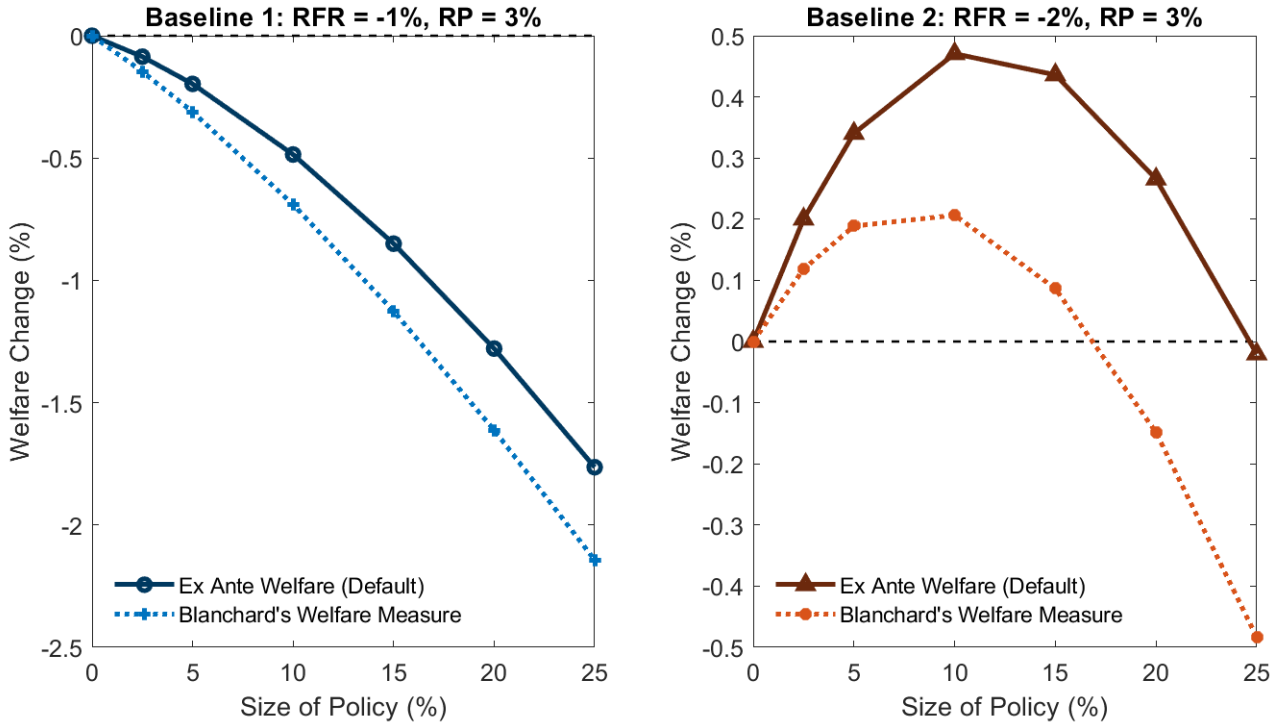


Figure 10: Comparing our EAU measure with Blanchard's measure. Impact of pay-go policy in the closed economy for both baselines.

deliver $c_1 = c_2 = 2$ (implying $U_1 = 2$) for sure. The welfare measure (20) implies indifference with respect to that policy. But why should uncertainty with respect to an agent's state at birth be evaluated risk-neutrally? In principle, any arbitrary degree of risk aversion with respect of the state of being born can be assumed. Blanchard's approach, which entails evaluating this risk based on a risk aversion coefficient of 1, simply amounts to³⁹

$$U_0 = 0.5 \log(U_1^A) + 0.5 \log(U_1^B). \quad (21)$$

Our alternative is to evaluate the uncertainty about the state in which an agent is born with the same degree of risk aversion with which the agent evaluates the uncertainty of old-age consumption. This leads to our ex-ante welfare measure

$$U_0 = (0.5(U_1^A)^{1-\gamma} + 0.5(U_1^B)^{1-\gamma})^{\frac{1}{1-\gamma}}, \quad (22)$$

which accords with EZ preferences as we now show.

³⁹Or in homogeneous form

$$U_0 = \exp(0.5(\log(U_1^A) + \log(U_1^B))).$$

C.3 Derivation of Ex-Ante Utility

We define Epstein-Zin utility recursively in a way that makes it homogeneous of degree one:

$$U_\tau^t = v^{-1} \left(v(c_\tau) + \tilde{\beta} v \left(u^{-1} \left(E_\tau \left[u \left(U_{\tau+1}^t \right) \right] \right) \right) \right) \text{ for } t \leq \tau < T, \text{ and } U_T^t = c_T,$$

with u capturing the attitude towards risk and v representing the attitude towards intertemporal substitution. We take $u(x) = x^{1-\gamma}/(1-\gamma)$ and $v(x) = \log(x)$. To determine ex-ante welfare of a generation born at $t > \tau$ we assume that agents born at t neither derive utility from consumption nor discount the future before birth – the same assumptions we would make in the time-separable case.⁴⁰ Consequently, the utility of a generation born at t , evaluated at, say, $\tau = t - 2$ becomes

$$\begin{aligned} U_{t-2}^t &= (v^{-1} \circ v \circ u^{-1}) \left(\mathbb{E}_{t-2} \left[(u \circ v^{-1} \circ v \circ u^{-1}) \left(\mathbb{E}_{t-1} \left[u(U_t^t) \right] \right) \right] \right) \\ &= u^{-1} \left(\mathbb{E}_{t-2} \left[\mathbb{E}_{t-1} \left[u(U_t^t) \right] \right] \right) = u^{-1} \left(\mathbb{E}_{t-2} \left[u(U_t^t) \right] \right). \end{aligned}$$

By iteration, the utility of a generation born at t evaluated $\tau = 0$ is

$$U_0^t = u^{-1} \left(\mathbb{E}_0 \left[u(U_t^t) \right] \right),$$

which amounts to (8) when using our specific utility function.

D Robustness and Sensitivity

We follow Blanchard (2019) and assume a high standard deviation of TFP and search for a risk-aversion parameter that matches the desired risk-premium. Table 5 provides sensitivity analysis with respect to this calibration method. The table's first three rows consider higher and lower values of σ , which implies lower and higher values of γ . Neither the B1 or B2 EAU results are much affected. The table's fourth row considers a calibration with a negatively skewed TFP shock. Specifically, we assume a disaster shock with a minus five standard deviation drop in TFP, which occurs each period with a 1 percent probability. All other TFP realizations are as before in the $\sigma = 0.2$ case, except that they are now slightly less likely and also slightly larger to keep average TFP constant. The long-run EAU results are slightly worse for this calibration. Interestingly, as Barro (2020)'s work and intuition suggest, the model now calibrates with much lower values of γ the model – 6.6 and 6.7 in B1 and B2, respectively.

⁴⁰Time-separable utility evaluated at time zero would be

$$U_0^t = E_0 \sum_{\tau=t}^{t+T} \tilde{\beta}^{\tau-t} u(c_\tau),$$

where a per-period utility function u captures both the attitude towards risk and intertemporal substitution.

TFP shocks	Baseline 1	Baseline 2
$\sigma = 0.15$	-1.27%	0.29%
$\sigma = 0.2$ (baseline)	-1.28%	0.27%
$\sigma = 0.3$	-1.30%	0.19%
$\sigma = 0.2 + \text{disaster}$	-1.40%	-0.01%

Table 5: Sensitivity of long-run EAU Changes from a 20% pay-go policy in the closed economy for changes in the distribution of TFP shocks

E Defined-Benefit and Constant-Debt Policies

All deficit policies aren't created equal. Blanchard's (defined-benefit) pay-go policy improves risk-sharing by maintaining transfers to the elderly regardless of the economy's state. This transfers risk to those who can best bear it – the young, thanks to their fixed endowment. This appendix considers two alternatives. First, a *defined-contribution* (DC) pay-go policy: $T_t = \kappa W_t$, for a fixed κ . Second, a policy that Blanchard (2019) considers as maintaining a constant level of debt, D . This policy entails $T_t = DR_t^f$ and will be called *constant debt* (CD) in what follows. Both alternatives are calibrated such that the long-run average transfer from the young to the old equals that under DB.

E.1 Alternative Policies in the Closed Economy

Table 6 shows, for the closed-economy case, that DC pay-go generates a long-run EAU loss under both B1 and B2. This impact can be decomposed as follows: DC produces smaller RNE losses and larger CRE gains compared to DB. But the LRE gains are zero under DC, whereas they are large under PG as the sure transfer in PG reduces the risk of old-age consumption. On balance, these factors make DC policy substantially worse than pay-go policy. As for constant debt policy, table 6 shows that it slightly dominates the PG policy in terms of its long-run EAU impact. This is, as table 6 reveals, entirely due to the improved CRE. CD improves risk-sharing between cohorts as it offers generations born with a bad shock a larger transfer when old than generations born with a good shock. This is because the constant debt scheme's transfer to the old in $t + 1$ is proportional to the interest rate R_{t+1}^f (accruing from t to $t + 1$), which is higher when there is a bad TFP shock in t .⁴¹ To summarize, the risk-sharing properties of deficit finance clearly matters to its EAU impact.

E.2 Alternative Policies in the Open Economy

Table 6 shows the long-run EAU effects for the two countries of defined-benefit, defined-contribution, constant debt, and two-way transfer policies. Compared to the closed-economy numbers (which are also provided in table 6), the results for all schemes are more favorable.

⁴¹The fact that low TFP coincides with high expected TFP growth and thus, all else equal, high interest rates is, in turn, a direct consequences of assuming i.i.d. realizations of the level of TFP. Therefore, the finding that CD dominates PG has to be taken with a grain of salt.

		Overall Effect			Decomposition								
Policy	Case	EAU change (in %)			RNE (in %)			LRE (in %)			CRE (in %)		
		Closed	Open		Cl.	Open		Cl.	Open		Cl.	Open	
		H	H	F	H	H	F	H	H	F	H	H	F
PG	B1	-1.3	0.04	-0.03	-3.2	-0.4	-0.5	1.6	0.2	0.3	0.4	0.2	0.2
DC	B1	-1.6	-0.90	0.08	-2.4	-1.7	-0.2	0.0	0.0	0.2	0.8	0.8	0.1
CD	B1	-0.9	0.01	0.01	-3.2	-0.4	-0.5	1.6	0.2	0.3	0.8	0.2	0.2
TT	B1	0.8	1.45	0.04	-0.9	0.4	-0.2	1.7	0.8	0.1	0.0	0.2	0.1
PG	B2	0.3	1.53	0.14	-2.7	0.8	-0.6	2.6	0.4	0.5	0.4	0.4	0.2
DC	B2	-1.0	-0.44	0.26	-2.0	-1.4	-0.1	0.0	-0.1	0.3	1.0	1.1	0.1
CD	B2	0.9	1.66	0.22	-2.7	0.8	-0.6	2.5	0.4	0.6	1.2	0.4	0.3
TT	B2	2.3	3.12	0.14	-0.8	0.9	-0.3	2.7	1.4	0.2	0.4	0.8	0.2

Table 6: Comparing long-run EAU impacts and their decomposition from different policies and calibrations in the two-country calibration (U.S. as the home country and the RoW as the foreign country). Closed-economy results reported for comparison.

This is due to the risk-neutral effect as the decomposition in table 6 reveals – and as was expected due to the dampening of crowding out owing to openness.

Comparing different deficit-finance schemes, we find that long-run EAU gains for the home country are much harder to achieve for a defined contribution scheme. This mirrors our results for the closed economy. However, the rest of the world now experiences long-run EAU gains under DC. The key to understanding this is that under DC the transfers are perfectly correlated with home-capital returns, making home-capital less attractive for home investors; consequently, the crowding out effects are much stronger for the home country than for the foreign country as can be seen from the RNE values in table 6. With crowding out being modest in the foreign country, risk-sharing effects dominate and the overall impact on foreigners is positive. When it comes to public debt, we observe that its home country impact now compares less favorably to the defined benefit scheme than in the closed economy. For the foreign country, however, CD is substantially better than DB.

F Calibrating the Two-Country Case

According to Penn World Tables data, the U.S. share of world GDP totaled 16.4% in 2017 – hence our assumption that the RoW is roughly six times larger than the U.S. We use the standard deviation of the growth rate of real GDP of the U.S. and RoW to approximate TFP risk. The standard deviation of the log difference in U.S. GDP growth is 1.95 log-percent, whereas that of the RoW is 2.50 log-percent. Therefore, we calibrate TFP of the RoW to be 1.25 times as risky as that of the U.S.. The same data produce our assumed cross-country TFP correlation of 0.22.

Finally, we introduce cross-country investment costs to match the 2019 cross-country asset positions reported by the Bureau of Economic Analysis. We adjust the model as follows.

The net cross-country investment return for home investors, $\hat{R}_{F,t+1}$, and the cross-country investment return for foreign investors, $\hat{R}_{H,t+1}$, are now given by

$$\hat{R}_{F,t+1} = R_{F,t+1} - \delta_F, \quad \hat{R}_{H,t+1} = R_{H,t+1} - \delta_H,$$

where δ_F and δ_H reference the costs to domestic and foreign residents of investing in capital abroad. We also introduce a bond investment cost, δ_{FB} , which foreigners face when investing in domestic bonds. $\hat{R}_{t+1}^f = R_{t+1}^f - \delta_{FB}$. Domestic investors face no cost when investing in domestic bonds and there is, by assumption, no separate foreign bond market. The budget constraints of the home and foreign old are

$$\begin{aligned} c_{o,H,t+1} &= k_{H,H,t+1}R_{H,t+1} + k_{H,F,t+1}\hat{R}_{F,t+1} + b_{H,t+1}R_{t+1}^f + T_{t+1} + TAC_{H,t+1}, \\ c_{o,F,t+1} &= k_{F,H,t+1}\hat{R}_{H,t+1} + k_{F,F,t+1}R_{F,t+1} + b_{F,t+1}\hat{R}_{t+1}^f + TAC_{F,t+1}, \end{aligned} \quad (23)$$

where $c_{o,H,t+1}$ is the consumption of the old in the home country, $k_{H,H,t}$ and $k_{H,F,t}$ are domestic and foreign capital investments made by the young at time t , $b_{H,t}$ is the time- t purchase of bonds by the young, T_{t+1} is the government transfer received by the domestic old at time $t+1$, and $TAC_{H,t}$ are the transaction costs, which we assume are lump-sum rebated. In equilibrium we have $TAC_{H,t} = \delta_H k_{H,F,t}$. Analogous formulations hold for foreign households.

References

- Abel, A. B., Mankiw, N. G., Summers, L. H., and Zeckhauser, R. J. (1989). Assessing dynamic efficiency: Theory and evidence. *The Review of Economic Studies*, 56(1):1–19.
- Abel, A. B. and Panageas, S. (2022). Optimal rollover of government bonds in a dynamically efficient economy.
- Aiyagari, S. R. and McGrattan, E. R. (1998). The optimum quantity of debt. *Journal of Monetary Economics*, 42(3):447–469.
- Aksoy, Y., Basso, H. S., Smith, R. P., and Grasl, T. (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, 11(1):193–222.
- Ball, L. and Mankiw, N. G. (2007). Intergenerational risk sharing in the spirit of Arrow, Debreu, and Rawls, with applications to social security design. *Journal of Political Economy*, 115(4):523–547.
- Ball, L. M. and Mankiw, N. G. (2021). Market power in neoclassical growth models. Working Paper 28538, National Bureau of Economic Research.
- Barbie, M., Hagedorn, M., and Kaul, A. (2007). On the interaction between risk sharing and capital accumulation in a stochastic OLG model with production. *Journal of Economic Theory*, 137(1):568–579.
- Barro, R. J. (2020). r minus g . Working Paper 28002, National Bureau of Economic Research.
- Blanchard, O. (2019). Public debt and low interest rates. *American Economic Review*, 109(4):1197–1229.
- Blanchard, O. (2023). *Fiscal policy under low interest rates*. MIT Press.
- Blanchard, O., Leandro, A., and Zettelmeyer, J. (2020). Revisiting the EU fiscal rules in an era of low interest rates. In *Rethinking the European Fiscal Framework Conference*.
- Brumm, J., Feng, X., Kotlikoff, L., and Kubler, F. (2022). Are deficits free? *Journal of Public Economics*, 208:104627.
- Brumm, J., Kotlikoff, L., and Kubler, F. (2020). Leveraging posterity’s prosperity? *AEA Papers and Proceedings*, 110:152–56.
- Brunnermeier, M. K., Merkel, S. A., and Sannikov, Y. (2020). Debt as safe asset: Mining the bubble. *Princeton University work in progress*.
- Cass, D. (1972). On capital overaccumulation in the aggregative, neoclassical model of economic growth: A complete characterization. *Journal of Economic Theory*, 4(2):200–223.

- Crafts, N. (2018). The productivity slowdown: Is it the ‘new normal’? *Oxford Review of Economic Policy*, 34(3):443–460.
- Davila, J., Hong, J. H., Krusell, P., and Ríos-Rull, J.-V. (2012). Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica*, 80(6):2431–2467.
- Dechert, W. D. (1982). Lagrange multipliers in infinite horizon discrete time optimal control models. *Journal of Mathematical Economics*, 9(3):285–302.
- Diamond, P. A. (1965). National debt in a neoclassical growth model. *The American Economic Review*, 55(5):1126–1150.
- Eggertsson, G. B., Mehrotra, N. R., and Robbins, J. A. (2019). A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics*, 11(1):1–48.
- Eggertsson, G. B., Mehrotra, N. R., Singh, S. R., and Summers, L. H. (2016). A contagious malady? open economy dimensions of secular stagnation. *IMF Economic Review*, 64(4):581–634.
- Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.
- Evans, R. W. (2020). Public debt, interest rates, and negative shocks. *AEA Papers and Proceedings*, 110:137–40.
- Evans, R. W., Kotlikoff, L. J., and Phillips, K. L. (2012). Game over: Simulating unsustainable fiscal policy. Working Paper 17917, National Bureau of Economic Research.
- Gordon, R. H. and Varian, H. R. (1988). Intergenerational risk sharing. *Journal of Public Economics*, 37(2):185–202.
- Gottardi, P. and Kubler, F. (2011). Social security and risk sharing. *Journal of Economic Theory*, 146(3):1078–1106.
- Green, J. and Kotikoff, L. J. (2008). On the general relativity of fiscal language. In Auerbach, A. J. and Shaviro, D. N., editors, *Institutional Foundations of Public Finance*, pages 241–256. Harvard University Press, Cambridge.
- Harenberg, D. and Ludwig, A. (2019). Idiosyncratic risk, aggregate risk, and the welfare effects of social security. *International Economic Review*, 60(2):661–692.
- Hasanhodzic, J. (2020). Simulating the Blanchard conjecture in a multiperiod life cycle model. *AEA Papers and Proceedings*, 110:149–51.

- Hasanhodzic, J. and Kotlikoff, L. J. (2013). Generational risk-is it a big deal?: Simulating an 80-period OLG model with aggregate shocks. Working Paper 19179, National Bureau of Economic Research.
- Hellwig, M. F. (2021). Safe assets, risky assets, and dynamic inefficiency in overlapping-generations economies. *MPI Collective Goods Discussion Paper*, (2021/10).
- Hubbard, R. G. and Judd, K. L. (1987). Social security and individual welfare: Precautionary saving, borrowing constraints, and the payroll tax. *The American Economic Review*, 77(4):630–646.
- İmrohoroglu, A., Imrohoroglu, S., and Joines, D. H. (1995). A life cycle analysis of social security. *Economic theory*, 6(1):83–114.
- Krueger, D. and Kubler, F. (2006). Pareto-improving social security reform when financial markets are incomplete!? *American Economic Review*, 96(3):737–755.
- Mian, A., Straub, L., and Sufi, A. (2021a). A goldilocks theory of fiscal policy. *NBER Working Paper*, (29351).
- Mian, A., Straub, L., and Sufi, A. (2021b). Indebted demand. *The Quarterly Journal of Economics*, 136(4):2243–2307.
- Miao, J. and Su, D. (2021). Fiscal and monetary policy interactions in a model with low interest rates. Working Paper, Boston University.
- Reis, R. (2021). The constraint on public debt when $r < g$ but $g < m$. Working paper, London School of Economics.
- Reis, R. (2022). The fiscal revenue from public borrowing. *Journal of Economic Perspectives*.
- Samuelson, P. A. (1958). An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy*, 66(6):467–482.
- Sergeyev, D. and Mehrotra, N. (2021). Debt sustainability in a low interest rate world. Discussion paper, Bocconi university.
- Sims, C. A. (2019). Optimal fiscal and monetary policy with distorting taxes. Working Paper, Princeton University.
- Summers, L. H. and Rachel, L. (2019). On falling neutral real rates, fiscal policy and the risk of secular stagnation. In *Brookings Papers on Economic Activity BPEA Conference Drafts*.
- Tirole, J. (1985). Asset bubbles and overlapping generations. *Econometrica: Journal of the Econometric Society*, 53(6):1499–1528.
- Zilcha, I. (1990). Dynamic efficiency in overlapping generations models with stochastic production. *Journal of Economic Theory*, 52(2):364–379.

Zilcha, I. (1991). Characterizing efficiency in stochastic overlapping generations models. *Journal of Economic Theory*, 55(1):1–16.