POSITIVE SCIENCE

Positive Science

Positive Science - Economics

Monetary Policy and the Intergenerational Incidence of Green Inflation

Matteo Bondesan
Independent Researcher

Publication Information

DOI: 10.65957/article.20

Positive Science - Economics, 2026, Article 20.

https://doi.org/10.65957/article.20

Citation: Article 20

Published: November 02, 2026

Publisher location: Strasbourg, France

OPEN ACCESS — Creative Commons Attribution (CC BY) 4.0

© 2026 by the author(s). Published by Positive Science SAS, Strasbourg, France.

This article is licensed under a Creative Commons Attribution 4.0 International License.

Monetary Policy and the Intergenerational Incidence of Green Inflation

Abstract

Carbon pricing is inflationary in the short run, and central banks must decide how hard to lean against this green inflation. I embed a monetary block–a reduced-form Phillips curve, an IS-type demand relation, and a Taylor-type policy rule–into a two-sector overlapping-generations transition economy in which carbon taxation generates a hump-shaped, cost-push inflation process. I derive an augmented green Taylor condition for the existence and sign-preserving well-posedness of the period-by-period equilibrium that is generally weaker than the classical unit threshold, and show, once the timing of the old's predetermined nominal bond is made explicit, that the monetary stance reallocates the burden of green inflation, at the point of impact, between the two cohorts alive at that date: a hawkish rule protects the old, whose nominal wealth would otherwise be eroded by realised inflation, but taxes the young through a larger output and wage contraction. A closed-form monetary-neutral coefficient equalises the two losses, independent of price stickiness and of the size of the cost-push shock. Because monetary policy alone cannot correct distribution without an efficiency cost, I show that, conditional on a lump-sum transfer and sufficient transfer revenue, pairing a stabilisation-only interest-rate rule with a fiscal transfer weakly–and, away from a knife-edge case, strictly–dominates using monetary policy for both purposes. Finally, I show that transitional monetary hawkishness raises the carbon stock at every finite horizon through a portfolio-substitution channel, with a gap that is exactly permanent only in the non-decaying limit and otherwise long-lived and only slowly decaying. This last result is conditional on two assumptions: that the cost-push path is unaffected by the monetary stance, and that brown-sector employment is a reduced-form decreasing function of the green-capital savings share, with no production structure behind it.

Keywords: Green inflation; Monetary policy; Taylor rule; OLG; Carbon tax; Intergenerational equity. **JEL Classification:** E31; E52; E63; H23; Q54.

1. Introduction

Carbon pricing raises the cost of fossil-fuel-intensive production before it lowers it. In the years immediately following the introduction or tightening of a carbon tax, brown-sector marginal cost rises, this cost passes through into the consumer price index, and the resulting green inflation–the transitory, policy-induced rise in the price level generated by the decarbonisation of the production structure–looks, from the vantage point of a central bank, indistinguishable from any other cost-push shock. Central banks then face a genuine choice: lean against this inflation with a conventional Taylor-type rule, or look through it as a supply-side disturbance that will fade once green productivity catches up. This choice is not merely a stabilisation problem. Because the young and the old hold different portfolios and earn income from different sources, the monetary response to green inflation redistributes resources across generations, and–as I show in this paper–it also affects how quickly the economy decarbonises. Understanding this three-way link between price stability, intergenerational equity, and the speed of the transition is the object of this paper.

I build on the structural model of Bondesan et al. (2026), henceforth BCB, which develops a two-sector overlapping-generations economy with a brown and a green production sector, a Pigouvian carbon tax, a convex climate-damage function, and lifecycle saving. That model shows that the energy transition generates a hump-shaped cost-push inflation process–inflationary while the tax phases in and brown prices adjust, deflationary once green learning-by-doing productivity gains dominate–and that this inflation burdens the old, who hold nominally fixed claims, more than the young, who are exposed mainly through wages. The framework is entirely real: there is no money, no nominal interest rate, and no central bank. This paper adds the missing nominal side. I import the model's cost-push inflation process as a primitive and layer on top of it a reduced-form New-Keynesian-style block–a Phillips curve, an IS-type demand relation, and a Taylor rule–together with a portfolio-choice margin linking the real bond return to the pace of green capital accumulation. This lets me ask, analytically, how the aggressiveness of the monetary response to green inflation reallocates the burden of the transition across generations, and whether it changes how fast the economy turns green. A word on the title: "intergenerational" refers here to the two cohorts alive at the same date–the old, who redeem a predetermined nominal claim, and the young, who earn wages–so the incidence results are a within-period comparison across cohorts, not a lifetime welfare comparison of a single generation (Section 3.6).

The paper connects four strands of the literature that have, to date, developed largely in parallel.

The first is the growing body of work embedding environmental policy into dynamic stochastic general equilibrium models with nominal rigidities. Annicchiarico and Di Dio (2015) study how emission-trading and carbon-tax rules interact with a New Keynesian business cycle, C. Fischer and Springborn (2011) and Heutel (2012) compare the cyclical properties of price- and quantity-based environmental instruments, and Diluiso et al. (2021) and McKibbin et al. (2020) study how monetary policy and financial regulation interact

with climate transition and stranded-asset risk. This literature is quantitative and typically representative-agent; it says little about intergenerational incidence, since there is, by construction, no heterogeneity across cohorts. Golosov et al. (2014) and Barrage (2020) embed optimal carbon taxation in general-equilibrium growth models with distortionary fiscal policy, which is close in spirit to the policy-assignment question I study in Section 4, but again without an overlapping-generations structure or a monetary authority.

The second is the overlapping-generations literature on environmental policy, which does have cohort heterogeneity but no nominal side. John and Pecchenino (1994) and Bovenberg and Heijdra (1998) study environmental taxation and intergenerational distribution in OLG economies, and Kotlikoff et al. (2021) quantify the generational winners and losers from carbon taxation in a large-scale life-cycle model. None of these papers has a central bank, a Phillips curve, or a nominal interest rate; the redistribution they study operates entirely through relative prices and public debt, not through the interaction of realised inflation with fixed nominal claims.

The third is the classical monetary-policy-rules literature, from Taylor (1993) to Clarida et al. (1999, 2000), Woodford (2003), and Galí (2015), which establishes the standard Taylor principle for determinacy and the trade-off between inflation and output-gap stabilisation. I import this apparatus but apply it to a cost-push process that is not a generic demand or productivity shock: it is a known, policy-engineered by-product of decarbonisation, with its own structural determinants (the speed of the carbon-tax phase-in and the pace of green learning-by-doing) and its own distributional incidence across generations that the standard representative-agent New Keynesian model cannot see.

The fourth strand studies how inflation itself redistributes resources between nominal debtors and creditors, independently of any carbon-tax context, and supplies the mechanism this paper's incidence result runs through. Doepke and Schneider (2006) document that inflation shocks transfer wealth from older, richer bondholding households to younger, middle-class debtors in the United States, and Auclert (2019) formalises the channels through which monetary policy redistributes consumption–an earnings-heterogeneity channel, a Fisher channel operating through unexpected inflation on nominal positions, and an interest-rate-exposure channel operating through the maturity structure of those positions–of which the Fisher and interest-rate-exposure channels are precisely the two forces at work in Section 3.2 below. On the green-inflation side specifically, Känzig (2023) identifies the distributional consequences of carbon-pricing shocks empirically; Olovsson and Vestin (2023) study the size and drivers of the price-level effects of decarbonisation in a New Keynesian two-sector model; and Ferrari and Nispi Landi (2024) study how central banks' unconventional asset-purchase policies can be tilted toward green assets. None of these papers combines the nominal-wealth-redistribution mechanism with an overlapping-generations structure and an explicit monetary-policy-versus-fiscal-transfer assignment question, which is what Sections 4.2–4.3 below do.

This paper's contribution is to bring these four lines of research together and de-

rive four key results. First, I show that the condition for the period-by-period system to have a unique, sign-preserving solution–the augmented green Taylor condition–is generally weaker than the classical unit Taylor threshold, because output-gap feedback in the policy rule also stabilises the system (Proposition 1); I return to a boundedness (stability) condition for the forward-looking inflation recursion once Section 5.1 introduces an expectational Phillips curve, without claiming a full rational-expectations determinacy result. Second, I show that, once the timing of the old's predetermined nominal bond is made explicit, the Taylor-rule inflation coefficient governs how the burden of green inflation splits, at the point of impact, between the old, who are protected from nominal-wealth erosion by a hawkish rule, and the young, who bear a larger output and wage contraction as a result; there is a unique, closed-form monetary-neutral coefficient that equalises the two losses, and it does not depend on the slope of the Phillips curve or on the size of the shock (Proposition 2 and Corollary 1). Third, since a lump-sum fiscal transfer can achieve the same equalisation at any monetary stance, I show that assigning monetary policy to stabilisation and fiscal transfers to equity–a green-inflation analogue of the classical assignment principle of Tinbergen (1952) and Mundell (1962)–weakly dominates, and away from a knife-edge case strictly dominates, using monetary policy for both purposes, conditional on the transfer being lump-sum and on sufficient transfer revenue (Proposition 3). Fourth, I show that, conditional on the cost-push path itself not responding to the monetary stance and on a reduced-form monotone mapping from the portfolio share to brown-sector activity, a more hawkish monetary stance raises, at every date, the real return that the young collect on nominal assets relative to a more dovish stance (under a condition on the slope of the cost-push path that the baseline calibration satisfies), slowing the reallocation toward the green sector and raising the carbon stock at every finite horizon–permanently only in the non-decaying limit, and otherwise with a gap that decays only slowly–so that transitory monetary hawkishness leaves a long-lived climate legacy (Proposition 4 and Corollary 2).

The remainder of the paper is organised as follows. Section 2 discusses the motivation, Section 3 sets out the model, while Section 4 states and proves the four main results and illustrates them numerically. Section 5 checks robustness to forward-looking price setting, interest-rate smoothing, vanishing portfolio sensitivity, and the calibration of the Phillips-curve slope and the wage-output elasticity, and discusses the model's limitations. Section 6 concludes with policy implications and directions for future work. Appendix A summarises notation and parameters, Appendix B collects the key equations, Appendix C reports the calibration and numerical results underlying the figures, and Appendix D contains all proofs.

2. Motivation and Stylised Facts

The model of Section 3 is motivated by two empirical patterns: that carbon pricing coincides with a visible, if imperfectly identified, cost-push pressure on consumer prices, and that the cost of external finance may matter for green investment, so that a change in the monetary stance could plausibly move the pace of decarbonisation. Neither pattern is new,

and I do not claim to identify either relationship causally from aggregate time series. Figure 1 is intended solely to motivate the coexistence of green-transition inflation, monetary tightening, and changing clean-energy investment: it documents that both patterns are visible in the data over the period in which the European carbon price and European policy rates moved the most. It should not be interpreted as empirical evidence for the causal portfolio mechanism developed in the model; its role is descriptive background, not a test.

Panel (a) plots the annual average EU Emissions Trading System (EU ETS) carbon price against euro area HICP energy inflation, 2017–2024. The carbon price rose from about EUR5 per tonne of CO2 in 2017 to a peak average of roughly EUR84 in 2023 before easing to about EUR65 in 2024, following the tightening of the Market Stability Reserve and the Fit-for-55 revision of the ETS directive (European Commission, Directorate-General for Economic and Financial Affairs, 2025; European Environment Agency, 2025). Energy HICP inflation in the euro area was close to zero through 2019, collapsed during the 2020 pandemic demand shock, then rose sharply–from an average of about 19 percent in 2021 to roughly 37 percent in 2022, before turning sharply negative in 2023–as base effects unwound (Eurostat, 2026). The co-movement between the two series over 2021–2022 is consistent with the cost-push mechanism of Assumption 1, though I emphasise that this was also the period of Russia's invasion of Ukraine and the associated gas-supply shock, which raised European energy prices for reasons unrelated to carbon policy; the aggregate correlation in panel (a) should therefore be read as suggestive of the pass-through channel the model formalises, not as evidence that isolates the carbon-price contribution from the concurrent, and considerably larger, gas-price shock.

Panel (b) plots the ECB deposit facility rate against global clean energy investment, 2018–2024. The policy rate was pinned at or below −0.40 percent from 2018 until mid-2022, then rose in the fastest hiking cycle in the euro's history to a terminal 4.00 percent in September 2023, before easing back to 3.00 percent by the end of 2024. Global clean energy investment grew throughout the period–reaching an estimated USD1.4 trillion in 2022, USD1.8 trillion in 2023, and USD2.0 trillion in 2024–but its growth rate decelerated sharply as policy rates peaked: from roughly 29 percent in 2023 to what the International Energy Agency itself describes as "moderate growth" of about 11 percent in 2024 (International Energy Agency, 2022, 2024, 2025). I do not read this deceleration as caused solely, or even primarily, by monetary policy: falling solar-module costs, the United States' Inflation Reduction Act, and industrial-policy support in China and the European Union all plausibly affected the pace of clean energy spending over the same years, and I have no way of separating these forces in aggregate data of this kind. The comparison is also weaker than panel (a) on three further counts: the two series differ in geographic coverage (a euro-area policy rate against global investment), the sample is very short, and part of the investment series is reconstructed from reported growth rates rather than read off single-year figures. The pattern is therefore compatible with, but does not test, the financing-cost channel embedded in the portfolio equation of Section 3.5, in which a higher real bond return crowds out

green capital at the margin; it motivates asking, as Section 4 does formally, how much of the cost of leaning against green inflation shows up not in the output gap but in the pace of the transition itself.

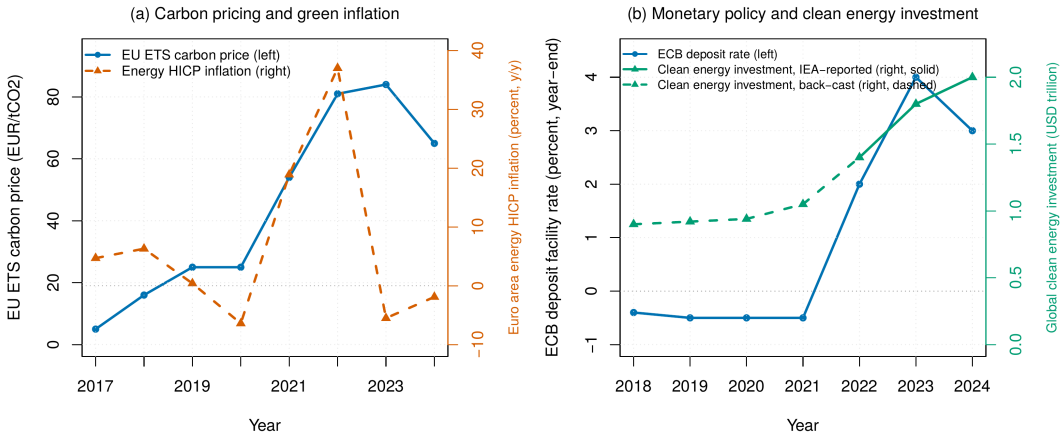


Figure 1. Motivating stylised facts

Source: Author's elaboration based on European Environment Agency (2025), European Commission, Directorate-General for Economic and Financial Affairs (2025) and Eurostat (2026) (panel a), and International Energy Agency (2022), International Energy Agency (2024) and International Energy Agency (2025) (panel b). *Note.* Panel (a): the EU ETS carbon price (EUR per tonne of CO2, annual average, left axis) against euro area HICP energy inflation (percent, year on year, right axis), 2017–2024. Panel (b): the ECB deposit facility rate (percent, year end, left axis) against global clean energy investment (USD trillion, right axis), 2018–2024. The 2018–2021 segment of the clean energy investment series in panel (b) is shown as a dashed line because it is back-cast from IEA-reported period-average growth rates rather than read directly off single-year IEA figures; the 2022–2024 segment is IEA-reported directly. Both panels document aggregate co-movement, not a causal estimate: the figure is descriptive motivation only and is not evidence for the causal portfolio mechanism of the model, and panel (b) compares a euro-area policy rate with global investment over a short sample.

These patterns are aggregate, short, and confounded by other shocks, and I treat them accordingly: as motivation for the mechanisms built into the model, not as an empirical test of it. Section 3 now sets out that model formally.

3. Model

I take as given the real block of the two-sector overlapping-generations transition economy of Bondesan et al. (2026) and add a nominal side. This section states the imported building block as a single assumption, then develops the new monetary and portfolio blocks in full.

3.1 Environment and the Imported Cost-Push Process

Time is discrete, $t = 0, 1, 2, \ldots$. At $t = 0$ an unanticipated carbon-tax reform moves the economy from a pre-existing, zero-tax, zero-inflation steady state onto a transition path with a positive and gradually phased-in carbon tax $\tau_t = \bar{\tau}(1 - e^{-\partial t})$, exactly as in the reference model. A continuum of two-period-lived agents is born each period; the young at t supply labour inelastically, earn a real wage w_t, consume C_t^1, and save; the old at t consume out of savings accumulated when young. Throughout, I normalise the inflation target to zero, $\pi^* = 0$, which simplifies notation without loss of generality for the qualitative results.

Length of a model period. One model period is one year. This is the only interpretation used to calibrate the model (Appendix C): the natural real rate $r^n = 0.02$, the carbon-stock decay rate $\mu = 0.05$ and the inflation peak at $\hat{T} = 5$ are annual magnitudes, the one-period nominal bond of Section 3.2 is a one-year contract, and the horizon of 30 periods in the figures is thirty years. The two-period-lived overlapping-generations structure is a stylised device for having, at every date, a working-age cohort that earns wages and an asset-holding cohort that redeems predetermined nominal claims; I do not interpret a life span as literally two years, and none of the analytical results depends on the length of the period. The period length matters only for the calibration, in particular for the Phillips-curve slope κ (Section 5.5).

Assumption 1 (Imported cost-push process)**.** The carbon-tax reform generates an exogenous–in particular, policy-invariant–hump-shaped cost-push component of consumer-price inflation, $\{\pi_t^{cp}\}_{t \geq 0}$, derived in the real block of the structural model from the pass-through of the tax into brown-sector prices net of green-sector learning-by-doing productivity gains. I approximate this process parametrically as

$$\pi_t^{cp} = \bar{\pi}^{cp} \frac{t}{\hat{T}} \exp\left(1 - \frac{t}{\hat{T}}\right), \qquad \bar{\pi}^{cp} > 0, \ \hat{T} > 0, \tag{1}$$

which is zero at $t = 0$, rises while the tax phase-in dominates, reaches its maximum $\bar{\pi}^{cp}$ at $t = \hat{T}$, and decays towards zero, without ever reaching it, as green productivity gains take over, consistent with the hump-shaped inflation path found in Bondesan et al. (2026).

Note that $\pi_t^{cp} > 0$ for every $t > 0$ and converges to zero only asymptotically; nothing below relies on a finite date after which it vanishes (this matters for Corollary 2).

The two parameters of Assumption 1 are structural objects of the model rather than free constants: $\hat{T}$ is pinned down by the relative speed of the carbon-tax phase-in, ∂, and the pace of green-sector learning-by-doing productivity growth, while the peak magnitude $\bar{\pi}^{cp}$ scales with the steady-state tax level $\bar{\tau}$ and the brown-sector cost pass-through elasticity; both are derived, together with the rest of the real block, in Bondesan et al. (2026). I take them as calibrated primitives here so as to keep the nominal side of the model self-contained.

Assumption 1 lets me treat the real side of the transition–sectoral reallocation, the carbon-tax schedule, and the climate-damage feedback–as already solved, and concentrate

the new analytical content of this paper on the monetary and portfolio margins that the reference model does not have. Everything from here on is new.

3.2 Nominal Assets, the Fisher Relation, and the Policy Rule

I first fix notation and timing precisely, since two of the paper's results turn on it. The sequence of events is as follows. Before $t = 0$ the economy is in the zero-tax, zero-inflation steady state, and all nominal claims are contracted under those conditions. At $t = 0$ the carbon-tax reform is announced unexpectedly; from $t = 0$ onward the whole transition path $\{\pi_t^{cp}\}_{t\geq 0}$ is known (Section 3.1). Within each period t: (a) the cost-push state π_t^{cp} is realised; (b) inflation π_t, the output gap $\hat{y}_t$ and the policy rate i_t are *jointly determined* as the solution of the Taylor rule (4) and the private-sector equilibrium conditions (5)–(6)–the rule is an equilibrium condition, not a literal sequence in which the central bank first observes $(\pi_t, \hat{y}_t)$ and then moves; (c) the young at t supply labour, earn w_t, consume C_t^1, and allocate savings between the nominal bond and green capital (Section 3.5) at the policy rate i_t; (d) the old at t–who purchased their nominal bond when young at $t-1$, at the gross rate $1+i_{t-1}$ fixed once and for all at that time–redeem it and consume C_t^2. Because the bond is a one-period, non-state-contingent nominal contract, its rate i_{t-1} is predetermined as of t: it was determined at $t-1$, jointly with $(\pi_{t-1}, \hat{y}_{t-1})$, and does not adjust when π_t–determined only afterwards–turns out to differ from π_{t-1}. In particular, the bond held by the first old cohort exposed to reform-induced inflation, that of $t = 1$, was bought at $t = 0$ at $i_0 = r^n$: because $\pi_0^{cp} = \hat{y}_0 = \pi_0 = 0$, the rule returns the pre-reform policy rate at $t = 0$ (see below), so this cohort holds predetermined nominal wealth whose realised real value is eroded by the reform-induced inflation of $t = 1$. This is the standard logic of predetermined nominal contracts (cf. S. Fischer, 1977): nothing about the old's bond return relies on any agent failing to foresee the transition path, only on the fact that the nominal return collected on a contract fixed one period earlier at the policy rate cannot respond to inflation determined afterwards.

Two real-rate objects therefore need to be kept distinct. The first is the *realised return on the old's predetermined bond*,

$$r_t^{old} \equiv i_{t-1} - \pi_t, \tag{2}$$

which uses the *previous* period's policy rate and enters the old's welfare loss (9) below (Section 3.6). Evaluated one period ahead it is exactly the return that the young at t collect on the bond they buy at t, $r_{t+1}^{old} = i_t - \pi_{t+1}$, and it is this object that drives the portfolio decision (7). The second is the *contemporaneous policy real rate*,

$$r_t = i_t - \pi_t, \tag{3}$$

which combines the *current* policy rate with current inflation. I interpret r_t as a reduced-form contemporaneous indicator of the monetary stance, used to summarise how tight policy is today in the static IS relation (6); it is not intended to coincide with the standard ex-ante

one-period real return on any asset, nor to be the ex-post return on any specific pre-existing contract. Section 4.2 uses r_t^{old} exclusively; equation (3) plays its role in the contemporaneous block of Section 4.1 and in the IS relation, and Section 4.4 uses r_{t+1}^{old}.

A central bank sets the nominal policy rate i_t according to a Taylor-type rule that responds to realised CPI inflation and to the output gap $\hat{y}_t$ (the log deviation of output from its flexible-price, or natural, level):

$$i_t = r_t^n + \phi_\pi \pi_t + \phi_y \hat{y}_t, \qquad \phi_\pi \geq 0,\ \phi_y \geq 0, \tag{4}$$

where r_t^n is the natural real rate consistent with the underlying real allocation, which I treat as a constant r^n for tractability. The coefficient ϕ_π is the central bank's hawkishness toward green inflation and is the paper's main policy parameter. Because $\pi_0^{cp} = \hat{y}_0 = \pi_0 = 0$ under Assumption 1 (the reform has not yet moved prices at the moment it occurs), $i_0 = r^n$ exactly, regardless of ϕ_π: the very first policy rate set after the reform is mechanically unaffected by it. This is what makes the incidence result of Section 4.2 exact at the point of impact.

3.3 Aggregate Supply: The Green Phillips Curve

Nominal prices are not fully flexible. I summarise the resulting short-run trade-off with a static, reduced-form Phillips curve in which realised inflation equals a demand-driven component, proportional to the output gap, plus the exogenous cost-push component imported from Assumption 1:

$$\pi_t = \kappa \hat{y}_t + \pi_t^{cp}, \qquad \kappa > 0. \tag{5}$$

The slope κ indexes the degree of price flexibility: as $\kappa \to \infty$ prices become fully flexible and all of π_t^{cp} passes directly into π_t regardless of monetary policy; for finite κ, monetary policy can trade off inflation against the output gap.

3.4 Aggregate Demand and the Natural Rate

Aggregate demand responds negatively to the gap between the real policy rate and the natural rate, in the manner of a standard IS relation:

$$\hat{y}_t = -\psi (r_t - r^n), \qquad \psi > 0. \tag{6}$$

Equations (3)–(6) form a static, reduced-form, New-Keynesian-style block–in the spirit of, but not equivalent to, the canonical forward-looking three-equation New Keynesian model–adapted to a two-period OLG environment in which each generation's decisions are made once and there is no operative expectational feedback beyond the perfect-foresight cost-push process itself. I return to a forward-looking version of equation (5) in Section 5.

3.5 Portfolio Choice and the Green-Capital Margin

The young allocate savings between the nominal bond of Section 3.2 and a claim on green-sector capital. I do not model this portfolio choice from micro-foundations; instead I posit a reduced-form portfolio rule in which the equilibrium green-capital share of aggregate savings, ρ_t^g, responds to the gap between the real return the young collect on the bond and the marginal product of green capital, MPK_t^g:

$$\rho_t^g \;=\; \bar{\rho} \;-\; \zeta\left(r_{t+1}^{old} - \mathrm{MPK}_t^g\right), \qquad \bar{\rho} \in (0,1),\ \zeta \geq 0, \tag{7}$$

where $r_{t+1}^{old} = i_t - \pi_{t+1}$ is the return on the one-period bond bought at t (equation (2) one period ahead), known at t under perfect foresight. The parameter ζ is the *return-spread sensitivity of green-capital allocation*: $\partial\rho_t^g/\partial(r_{t+1}^{old} - \mathrm{MPK}_t^g) = -\zeta$, so a larger ζ makes the green-capital share more responsive to the return differential, and $\zeta = 0$ makes it entirely unresponsive. One way to rationalise (7) is a quadratic cost $\frac{c}{2}(\rho - \bar{\rho})^2$ of moving savings away from the baseline allocation–liquidity preference, collateral constraints, or the bank-lending relationships that intermediate green investment: maximising the return net of this cost gives (7) with $\zeta = 1/c$, so that $\zeta \to 0$ is an infinitely rigid portfolio and perfect substitutability would correspond to $\zeta \to \infty$. I use this only as an interpretation; the rule is reduced-form and ζ is not disciplined by independent evidence.

I evaluate equation (7) at $\mathrm{MPK}_t^g = r^n$, its pre-reform equilibrium value, so that $\rho_t^g = \bar{\rho} - \zeta(r_{t+1}^{old} - r^n)$: whenever the return the young collect on the bond rises above the natural rate, savings are reallocated away from green capital and toward the nominal bond. Equation (7) is linear and unbounded, so the restriction $\rho_t^g \in [0,1]$ is not automatic; it is derived, as a closed-form parameter condition, in Section 4.4.

Assumption 2 (Reduced-form environmental mapping)**.** The brown-sector employment share is a strictly decreasing, Lipschitz-continuous function of the green-capital savings share, $L_t^b = f(\rho_t^g)$, with baseline $f(\rho) = 1 - \rho$ (normalising the total factor supply to one). Emissions are $e_t = \varepsilon L_t^b$, and the carbon stock accumulates as

$$Z_t \;=\; (1-\mu)\,Z_{t-1} \;+\; \varepsilon\, L_{t-1}^b, \qquad \mu \in (0,1), \tag{8}$$

from a common initial condition Z_0 ($Z_0 = 0$ in the de-meaned units of the numerical illustration), exactly as in the BCB model, with μ the atmospheric decay rate of the carbon stock.

Assumption 2 is a reduced-form normalisation introduced to close the climate channel parsimoniously. It is not derived from a sectoral production problem: a financial portfolio share and a labour-allocation share are conceptually distinct objects, and the mapping only captures the stylised idea that a larger allocation of savings toward green capital goes with a smaller brown-sector activity share. Proposition 4 and Corollary 2 are therefore conditional on this monotone mapping (and, by Assumption 1, on a policy-invariant cost-push

path); the numerical illustrations use the baseline $f(\rho) = 1 - \rho$. Emissions of period $t-1$ enter the stock at t, which is the timing used in the proofs and in the code.

3.6 Fiscal Transfers and Cohort Welfare

The government levies the carbon tax and can rebate a share $\xi \in [0,1]$ of the revenue T_t to the contemporaneous old, with the remainder $(1-\xi)T_t$ going to the young, exactly as in the BCB model's revenue-neutral reform. I evaluate the first-order welfare consequences of the transition for the two cohorts alive at any date t using linear approximations around the pre-reform steady state. The old's loss from nominal-wealth erosion uses the realised bond return r_t^{old} of equation (2):

$$\begin{aligned}\mathcal{L}_t^O &= u'(C^2)\,\bar{b}\left(r^n - r_t^{old}\right) \;-\; u'(C^2)\,\xi\,\frac{T_t}{P_t} \\ &= u'(C^2)\,\bar{b}\left(\pi_t - \phi_\pi \pi_{t-1} - \phi_y \hat{y}_{t-1}\right) \;-\; u'(C^2)\,\xi\,\frac{T_t}{P_t},\end{aligned} \tag{9}$$

where $\bar{b}$ is the old's predetermined real bond wealth, $u'(C^2)$ their marginal utility of old-age consumption, and the second equality substitutes $i_{t-1} = r^n + \phi_\pi \pi_{t-1} + \phi_y \hat{y}_{t-1}$ from the Taylor rule evaluated at $t-1$. Because $\pi_0 = \hat{y}_0 = 0$ (Section 3.2), this collapses to $\mathcal{L}_1^O = \bar{b}\,\pi_1$ at $t=1$, the form used in Proposition 2 below; Remark 2 gives the general-t expression. The young's loss from the output and real-wage contraction is

$$\mathcal{L}_t^Y \;=\; -u'(C^1)\,\bar{w}\,\nu\,\hat{y}_t \;-\; u'(C^1)\,(1-\xi)\,\frac{T_t}{P_t}, \tag{10}$$

where $\bar{w}$ is the flexible-price steady-state real wage, $\nu > 0$ is the elasticity linking the real wage to the output gap ($w_t = \bar{w}(1+\nu\hat{y}_t)$), and $u'(C^1)$ is the young's marginal utility. Throughout Sections 3–4 I normalise $u'(C^1) = u'(C^2) = 1$ and set $\xi = 0$ unless stated otherwise, so that $\mathcal{L}_t^O$ and $\mathcal{L}_t^Y$ isolate the pure monetary-incidence channel; Section 4.3 reintroduces ξ.

Two clarifications follow directly from the timing of Section 3.2 and are worth making explicit. First, $\mathcal{L}_t^O$ and $\mathcal{L}_t^Y$ are, by construction, a *within-period* comparison: they measure the first-order welfare consequence experienced, at the same date t, by the two different cohorts who happen to be alive then–the old, exposed through a predetermined nominal claim, and the young, exposed through the current labour market–not a lifetime welfare comparison of a single generation across its two periods of life. Second, and precisely because of this, the young's own portfolio choice at t (Section 3.5) does not enter $\mathcal{L}_t^Y$: any real-return consequence of that choice materialises only when this cohort is old, at $t+1$, and shows up–for that cohort, at that later date–in $\mathcal{L}_{t+1}^O$ through r_{t+1}^{old} once its bond position is realised–the same return that drives the portfolio rule (7). Reading Proposition 2 and Corollary 1 as a lifetime-welfare comparison, rather than as the within-period transfer they are designed to capture, would be a misreading of what the closed form computes; Section 5.5 returns to what a lifetime-welfare extension would require.

3.7 Equilibrium Definition

Definition 1 (Monetary transition equilibrium)**.** Given the cost-push process $\{\pi_t^{cp}\}_{t\geq 0}$ of Assumption 1 and policy parameters (ϕ_π, ϕ_y, ξ), a monetary transition equilibrium is a sequence $\{\pi_t, \hat{y}_t, i_t, r_t, r_t^{old}, \rho_t^g, Z_t\}_{t\geq 0}$ satisfying the Fisher relation (3), the old's realised-return identity (2), the Taylor rule (4), the Phillips curve (5), the IS relation (6), the portfolio equation (7), and the carbon-stock law of motion (8), each period, together with the cohort welfare losses (9)–(10).

Because equations (3)–(6) are static within each period given π_t^{cp} and r^n, the equilibrium sequence $\{\pi_t, \hat{y}_t\}$ solves independently period by period; the only genuinely dynamic element is the carbon stock (8), which accumulates the period-by-period outcomes. I solve this system in closed form in Section 4.

4. Results

4.1 Existence and Well-Posedness: The Augmented Green Taylor Condition

Solving equations (3), (4), (5), and (6) simultaneously for $(\pi_t, \hat{y}_t)$ given π_t^{cp} yields a single scalar that governs existence and uniqueness of the period-by-period equilibrium. Because this baseline system is static, solved separately in each period with no term linking π_t to π_{t+1}, the condition below is an invertibility and sign-preservation condition for a 2×2 linear system, not a statement about the multiplicity of rational-expectations equilibria that the word *determinacy* refers to in the New Keynesian literature; I keep the language of a 'Taylor principle' only because condition (13) nests the classical threshold as a limiting case, and I return to a boundedness (stability) condition for the forward-looking inflation recursion in Section 5.1 once the Phillips curve is made forward-looking, without claiming a full rational-expectations determinacy result for this reduced-form block.

Proposition 1 (Augmented green Taylor condition: existence and well-posedness)**.** Define

$$\Delta(\phi_\pi) \;\equiv\; 1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1). \tag{11}$$

For any π_t^{cp}, the system (3), (4), (5), (6) has a unique, finite solution

$$\pi_t \;=\; \frac{(1+\psi\phi_y)\,\pi_t^{cp}}{\Delta(\phi_\pi)}, \qquad \hat{y}_t \;=\; \frac{\psi(1-\phi_\pi)\,\pi_t^{cp}}{\Delta(\phi_\pi)}, \tag{12}$$

if and only if $\Delta(\phi_\pi) \neq 0$; the solution is well-posed, in the sense that π_t has the same sign as π_t^{cp}, if and only if

$$\phi_\pi \;>\; \phi_\pi^{\min} \;\equiv\; 1 - \frac{1+\psi\phi_y}{\kappa\psi}. \tag{13}$$

The proof is in Appendix D. Condition (13) is the augmented green Taylor condition: as the Phillips curve becomes vertical ($\kappa \to \infty$, i.e., prices become fully flexible), $\phi_\pi^{\min} \to 1$ and the classical Taylor threshold is recovered exactly, though–absent forward-

looking expectations in this baseline–recovering the threshold's numerical value does not by itself recover the rational-expectations determinacy property that threshold is usually invoked for (Section 5.1 studies a stability condition for a forward-looking version). For any finite κ and any $\phi_y > 0$, $\phi_\pi^{\min} < 1$, so the requirement is strictly weaker than the textbook threshold: output-gap feedback in the policy rule contributes to stabilising the system on top of the direct inflation-feedback channel, so a central bank need not react as aggressively to inflation as the flexible-price benchmark suggests once it also responds to the output gap.

Remark 1. Equation (12) shows that realised inflation is a constant fraction $(1 + \psi\phi_y)/\Delta(\phi_\pi)$ of the underlying cost-push process, and this fraction is strictly decreasing in ϕ_π whenever condition (13) holds and $\phi_\pi > 1$: a more hawkish central bank mechanically absorbs a larger share of green inflation into a negative output gap rather than into realised prices.

4.2 The Monetary Incidence of Green Inflation

I now use Proposition 1, together with the timing distinction of Section 3.2, to characterise how the monetary stance reallocates the burden of the transition between the two cohorts alive at a given date. I evaluate the cohort losses (9)–(10) (with $\xi = 0$) at $t^* = 1$, the first date at which green inflation is realised against a bond bought at the pre-reform policy rate $i_0 = r^n$. This choice is not merely expositional: because $i_0 = r^n$ exactly (Section 3.2), the old's loss at $t = 1$ takes the simple, one-term form $\mathcal{L}_1^O = \bar{b}\,\pi_1$, with no approximation involved. At later dates $t > 1$, equation (9) involves both π_t^{cp} and π_{t-1}^{cp}, since the bond's predetermined rate reflects the *previous* period's, not the current period's, cost-push realisation; Remark 2 below gives this general expression and discusses what it implies away from the point of impact. Proposition 2 thus characterises the incidence on the two cohorts exposed at the point of impact following the unanticipated reform–the first old cohort, holding a claim bought at the pre-reform policy rate, and the young of the same date; it is not a statement that applies mechanically to every cohort throughout the transition.

Proposition 2 (Monetary incidence)**.** Let $t^* = 1$, $\pi_*^{cp} \equiv \pi_1^{cp} > 0$, and suppose $\phi_\pi > \phi_\pi^{\min}$. Define

$$\mathcal{L}^O(\phi_\pi) \equiv \bar{b}\,\pi_{t^*}(\phi_\pi), \qquad \mathcal{L}^Y(\phi_\pi) \equiv -\bar{w}\,\nu\,\hat{y}_{t^*}(\phi_\pi). \tag{14}$$

Then, for $\phi_\pi > 1$: (i) $\mathcal{L}^O(\phi_\pi)$ is strictly decreasing, with $\mathcal{L}^O(1) = \bar{b}\,(1 + \psi\phi_y)\pi_*^{cp}/(1 + \psi\phi_y) > 0$ and $\lim_{\phi_\pi\to\infty}\mathcal{L}^O(\phi_\pi) = 0$; (ii) $\mathcal{L}^Y(\phi_\pi)$ is strictly increasing, with $\mathcal{L}^Y(1) = 0$ and $\lim_{\phi_\pi\to\infty}\mathcal{L}^Y(\phi_\pi) = \bar{w}\nu\pi_*^{cp}/\kappa$; (iii) there exists a unique $\phi_\pi^* > 1$ such that $\mathcal{L}^O(\phi_\pi^*) = \mathcal{L}^Y(\phi_\pi^*)$, with the old bearing the larger loss for $\phi_\pi < \phi_\pi^*$ and the young bearing the larger loss for $\phi_\pi > \phi_\pi^*$.

Corollary 1 (Closed-form monetary-neutral coefficient)**.** The monetary-neutral

coefficient of Proposition 2 has the closed form

$$\phi_\pi^* \;=\; 1 + \frac{\bar{b}\,(1+\psi\phi_y)}{\bar{w}\,\nu\,\psi}, \tag{15}$$

which is independent of the Phillips-curve slope κ and of the magnitude of the cost-push shock π_*^{cp}.

The proofs are in Appendix D. Corollary 1 is a genuinely convenient result for policy design: the monetary stance that equalises the first-order burden of green inflation across generations can be computed from three ratios–the old's real bond wealth relative to the young's wage bill, the IS-curve interest sensitivity, and the output-gap coefficient already embedded in the policy rule–without needing to know how sticky prices are or how large the carbon-tax shock will turn out to be.

Remark 2 (General-t incidence)**.** At a general date $t \geq 1$, substituting $i_{t-1} = r^n + \phi_\pi \pi_{t-1} + \phi_y \hat{y}_{t-1}$ and the equilibrium expressions (12) for $(\pi_{t-1}, \hat{y}_{t-1})$ into equation (9) gives

$$\mathcal{L}_s^O(\phi_\pi) \;=\; \frac{\bar{b}}{\Delta(\phi_\pi)}\Big[(1+\psi\phi_y)\,\pi_t^{cp} - (\phi_\pi + \psi\phi_y)\,\pi_{t-1}^{cp}\Big], \tag{16}$$

which reduces to $\bar{b}\,(1+\psi\phi_y)\pi_1^{cp}/\Delta(\phi_\pi)$ at $t = 1$ (since $\pi_0^{cp} = 0$), recovering Proposition 2 exactly. For $t > 1$ the crossing point that equalises $\mathcal{L}_s^O(\phi_\pi)$ and $\mathcal{L}_s^Y(\phi_\pi)$ is no longer given by the date-invariant closed form (15), because $\mathcal{L}_s^O$ now depends on the *change* in the cost-push path around t, $\pi_t^{cp} - \pi_{t-1}^{cp}$, and not only on its level–an earlier version of this remark claimed date-invariance at every t, which this correction supersedes. Under the baseline calibration of Appendix C, the crossing point falls rapidly as t moves away from the point of impact: $\phi_\pi^*(1) = 4.00$, and $\phi_\pi^*(2), \ldots, \phi_\pi^*(5) \approx 1.53,\ 1.21,\ 1.09,\ 1.02$ as t approaches the peak $\hat{T} = 5$, where the cost-push path is locally flat. Beyond the peak, once π_t^{cp} is falling, $\mathcal{L}_s^O(\phi_\pi)$ is negative for every $\phi_\pi > 1$ in this calibration–the old *gain* at every hawkish stance, because the predetermined rate set during the still-rising phase of the hump overcompensates them relative to the now-lower realised inflation–so no crossing exists in the hawkish region and the young bear the larger loss throughout. The clean, single-crossing, date-invariant story of Proposition 2 and Corollary 1 is therefore a feature of the point of impact specifically; away from it, incidence is governed by the local slope, not the level, of the cost-push path, and the qualitative direction of the effect can reverse once the process turns down. Two further caveats apply away from the point of impact. First, later cohorts have seen the reform and adjust their portfolios through equation (7); because equation (9) is a first-order approximation around the pre-reform steady state, holding $\bar{b}$ fixed is innocuous only to first order (a portfolio-induced change in $\bar{b}$ multiplies a return deviation that is itself first order), which is a further reason to read the cleanest incidence result as one about the initially exposed cohorts. Second, the comparison remains a within-period one (Section 3.6).

4.3 Policy Assignment: Stabilisation versus Equity

Proposition 2 shows that monetary policy can equalise the cohort burden of green inflation, but only by pushing the output gap away from zero, which is costly in terms of aggregate stabilisation. I now show that this cost is avoidable, because the fiscal instrument ξ of Section 3.6 can achieve the same equalisation without any output cost.

Proposition 3 (Policy assignment)**.** Let $\tilde{\mathcal{L}}_s^O(\phi_\pi, \xi) = \mathcal{L}_s^O - \xi\, T_t/P_t$ and $\tilde{\mathcal{L}}_s^Y(\phi_\pi, \xi) = \mathcal{L}_s^Y - (1-\xi)T_t/P_t$ denote the net cohort losses inclusive of the revenue-neutral transfer. Then: (i) for any $\phi_\pi > \phi_\pi^{\min}$ satisfying $\Delta(\phi_\pi) > 0$, there exists a unique transfer share $\xi^{**}(\phi_\pi)$ such that $\tilde{\mathcal{L}}_s^O = \tilde{\mathcal{L}}_s^Y$, and it is interior, $\xi^{**}(\phi_\pi) \in (0,1)$, if and only if the available transfer revenue exceeds the incidence gap, $T_t/P_t > |\mathcal{L}_s^O - \mathcal{L}_s^Y|$; (ii) consequently, whenever this feasibility condition holds, for a policymaker minimising a loss function of the form $\mathcal{W}(\phi_\pi, \xi) = \frac{1}{2}\left(\pi_t^2 + \omega_y \hat{y}_t^2\right) + \omega_d\left(\tilde{\mathcal{L}}_s^O - \tilde{\mathcal{L}}_s^Y\right)^2$ with $\omega_y, \omega_d > 0$, the jointly optimal policy sets $\xi = \xi^{**}(\phi_\pi)$ for any ϕ_π, reducing the distributional term to zero, and then chooses ϕ_π to minimise the stabilisation term $\frac{1}{2}(\pi_t^2 + \omega_y \hat{y}_t^2)$ alone, independently of ω_d.

The proof is in Appendix D. Part (ii) is a green-inflation instance of the classical policy-assignment principle of Tinbergen (1952) and Mundell (1962): with two independent targets (price/output stability and intergenerational equity) and two independent instruments (the Taylor-rule coefficient and the fiscal transfer share), the efficient assignment pairs each instrument with the target it affects at lowest cost. Using ϕ_π alone to achieve equity, as in Proposition 2, is feasible but weakly dominated by the joint policy of Proposition 3, and strictly dominated away from the knife-edge case in which the stabilisation-optimal ϕ_π already equalises the two cohort losses: outside that case, using ϕ_π alone forces the central bank away from its stabilisation optimum and imposes an avoidable output cost.

Read precisely, part (i) says: conditional on sufficient available transfer revenue and lump-sum redistribution, the fiscal authority can choose a transfer share that equalises the contemporaneous losses of the two cohorts. This is weaker than saying that redistribution can always neutralise intergenerational incidence. The feasibility condition is also quantitatively informative: at the point of impact under the baseline calibration the incidence gap $|\mathcal{L}_1^O - \mathcal{L}_1^Y|$ equals 0.0205, 0.0111 and 0.0007 (model units) under the dovish, benchmark and hawkish rules, so the revenue needed is largest precisely when monetary policy does the least distributional work.

The result is conditional on three features of the policymaker's problem in Proposition 3: the loss function is additively separable in a stabilisation term and a squared distributional term; the transfer ξ is lump-sum, so that equalising cohort losses through ξ carries no efficiency cost of its own; and available transfer revenue is sufficient, as in the feasibility condition of part (i). Under a separable, no-cost objective, the assignment result follows almost immediately once a transfer capable of hitting the distributional target exactly is available; Proposition 3 should therefore be read as formalising the Tinbergen logic in this specific setting, with the substantive economic content residing in the lump-sum-transfer assump-

tion, which I now relax. If instead the transfer is distortionary–for instance, if raising and administering $\xi T_t/P_t$ costs $\frac{\chi}{2}(\xi T_t/P_t)^2$ in foregone output for some $\chi > 0$–the separability that drives Proposition 3 breaks down: the policymaker's problem becomes

$$\min_{\phi_\pi,\xi} \ \tfrac{1}{2}\left(\pi_t^2 + \omega_y \hat{y}_t^2\right) + \omega_d\left(\tilde{\mathcal{L}}_t^O - \tilde{\mathcal{L}}_t^Y\right)^2 + \tfrac{\chi}{2}\left(\xi\,\frac{T_t}{P_t}\right)^2, \tag{17}$$

in which the marginal cost of using ξ for redistribution is increasing in ξ, so the jointly optimal policy generally sets $\xi < \xi^{**}(\phi_\pi)$ and lets ϕ_π absorb part of the distributional burden even away from the stabilisation optimum: the clean assignment becomes a genuine trade-off, governed by the ratio of χ to ω_d, rather than a strict separation. I leave the full characterisation of this distortionary-transfer case–including the resulting joint first-order conditions in (ϕ_π, ξ)–for future work, and read Proposition 3 as establishing the *direction* of the assignment, not its magnitude, once $\chi > 0$. I discuss the lump-sum idealisation further in Section 5.

4.4 Monetary Policy and the Speed of the Green Transition

The results so far concern a single date. I now consider a monetary stance ϕ_π applied at every date of the transition, and ask how it affects the pace of decarbonisation through the portfolio channel of equation (7). Two conditions are maintained throughout this subsection, and Proposition 4 and Corollary 2 are conditional on both. First, the cost-push path $\{\pi_t^{cp}\}$ of Assumption 1 is held *fixed* across the two monetary regimes compared below: the results are statements about the climate consequences of the monetary stance *conditional on* a common, policy-invariant cost-push path, not an unconditional climate effect of hawkishness in a world where the pace of decarbonisation could itself alter the inflation process (for instance, through the damage function or through green learning-by-doing in the BCB model's real block). Second, the link from portfolios to emissions is the reduced-form monotone mapping $L_t^b = f(\rho_t^g)$ of Assumption 2, with no production structure behind it. Section 5.5 discusses the feedback and the sectoral structure that these two conditions omit.

The return the young collect. Under the timing of Section 3.2, the young at t collect $r_{t+1}^{old} = i_t - \pi_{t+1}$ on the bond they buy at t. Using the Taylor rule (4) and the equilibrium expressions (12) for $(\pi_t, \hat{y}_t)$ and for π_{t+1}, this return exceeds the natural rate by

$$r_{t+1}^{old} - r^n \;=\; \frac{(\phi_\pi + \psi\phi_y)\,\pi_t^{cp} - (1 + \psi\phi_y)\,\pi_{t+1}^{cp}}{\Delta(\phi_\pi)}, \tag{18}$$

and its slope in the policy coefficient is

$$\frac{\partial}{\partial \phi_\pi}\left(r_{t+1}^{old} - r^n\right) \;=\; \frac{(1 + \psi\phi_y)\,S_t}{\Delta(\phi_\pi)^2}, \qquad S_t \;\equiv\; (1 - \kappa\psi)\,\pi_t^{cp} + \kappa\psi\,\pi_{t+1}^{cp}. \tag{19}$$

Unlike the contemporaneous stance indicator $r_t - r^n = (\phi_\pi - 1)\pi_t^{cp}/\Delta(\phi_\pi)$, whose slope is

positive by algebra alone, the sign of (19) depends on $\kappa\psi$ and on the local slope of the cost-push path. A higher ϕ_π lowers next period's inflation, which raises the return; but when $\kappa\psi > 1$ it also lowers the equilibrium policy rate i_t, because inflation falls by more than the rule's coefficient rises; the net sign is that of S_t. At $t = 0$, $S_0 = \kappa\psi\,\pi_1^{cp} > 0$. For $t \geq 1$, $S_t = \pi_t^{cp}[\,1 - \kappa\psi + \kappa\psi\, q_t\,]$ with $q_t \equiv \pi_{t+1}^{cp}/\pi_t^{cp}$, so $S_t > 0$ if $\kappa\psi \leq 1$, and, for $\kappa\psi > 1$, if and only if $q_t > 1 - 1/(\kappa\psi)$. Under Assumption 1, $q_t = (1 + 1/t)\,e^{-1/\hat{T}}$ is strictly decreasing in t and converges to $e^{-1/\hat{T}}$, so $S_t > 0$ at every date $t \geq 0$ if and only if

$$\kappa\psi \;\leq\; \frac{1}{1 - e^{-1/\hat{T}}}. \tag{20}$$

The right-hand side equals 5.52 at $\hat{T} = 5$, so the baseline $\kappa\psi = 2$ satisfies (20) with a wide margin; the monotonicity of the return in ϕ_π is nevertheless a calibration-dependent property rather than an algebraic one. Note also that $r_{t+1}^{old} - r^n$ itself need not be positive: it is negative at $t = 0$ for every ϕ_π (there $i_0 = r^n$ while $\pi_1 > 0$), and, for moderate ϕ_π, during the rising phase of the hump (at the baseline, for $t = 0, \ldots, 3$ under the dovish rule but only at $t = 0$ under the hawkish rule). What matters for the comparative statics is the ordering across stances.

Interiority. Equation (7) is linear and unbounded, so $\rho_t^g \in (0, 1)$ is a restriction on parameters, not an implication of the rule. Given the closed form (18), the share is interior at every date if $\zeta \sup_t |r_{t+1}^{old} - r^n| < \min\{\bar{\rho}, 1 - \bar{\rho}\}$, and, because $|A - B| \leq \max\{A, B\}$ for the two non-negative terms in the numerator of (18), a sufficient closed-form condition is

$$\zeta\,\frac{\max\{\phi_\pi, 1\} + \psi\phi_y}{\Delta(\phi_\pi)}\;\bar{\pi}^{cp} \;<\; \min\{\bar{\rho},\, 1 - \bar{\rho}\}. \tag{21}$$

The comparative statics below are stated for parameter values, and for the two stances compared, such that the implied green-capital share remains interior, $\rho_t^g \in (0, 1)$. At the baseline the exact left-hand-side quantity $\zeta \sup_t |r_{t+1}^{old} - r^n|$ is 0.0083 under both the dovish and the hawkish rule (the sufficient bound (21) gives 0.019 and 0.012) against a limit of 0.5, so the shares are interior with a very wide margin (any ζ up to about 24 would keep them so); the same holds in every sensitivity exercise of Appendix C (Table 9).

Proposition 4 (Monetary policy and the transition speed)**.** Maintain Assumptions 1 and 2, and consider two stances $\phi_\pi^H > \phi_\pi^D > \phi_\pi^{\min}$ ("hawkish" and "dovish") for which the green-capital share is interior, condition (21). Then: (i) $\partial(r_{t+1}^{old} - r^n)/\partial\phi_\pi$ has the sign of S_t; $S_0 > 0$ always, and $S_t > 0$ for every $t \geq 0$ if and only if condition (20) holds; (ii) under (20), $r_{t+1}^{old,H} > r_{t+1}^{old,D}$ for every $t \geq 0$; (iii) consequently, if $\zeta > 0$, the induced green-capital-share paths satisfy $\rho_t^{g,H} < \rho_t^{g,D}$, and brown-sector employment $L_t^{b,H} > L_t^{b,D}$, for every $t \geq 0$; (iv) the resulting carbon-stock paths satisfy $Z_t^H \geq Z_t^D$ for every $t \geq 0$, with equality at $t = 0$ and, if $\zeta > 0$, strict inequality for every $t \geq 1$.

Corollary 2 (Long-lived climate legacy of transitory hawkishness)**.** Under the con-

ditions of Proposition 4 and $\zeta > 0$, let $e_t \equiv \varepsilon\,[f(\rho_t^{g,H}) - f(\rho_t^{g,D})] > 0$ denote the emissions differential and $D_t \equiv Z_t^H - Z_t^D$ the carbon-stock gap. As the transition-induced inflation shock converges to zero, the emissions differential converges to zero, $e_t \to 0$, and is summable, $\sum_{t\geq 0} e_t < \infty$. Consequently: (a) if $\mu \in (0,1)$, the gap converges to zero (geometrically), $D_t \to 0$ as $t \to \infty$, while remaining strictly positive at every $t \geq 1$; (b) if $\mu = 0$, the gap is strictly increasing and converges to a strictly positive constant,

$$\lim_{t\to\infty} \left(Z_t^H - Z_t^D\right) \;=\; \varepsilon \sum_{s=0}^{\infty} \left[f(\rho_s^{g,H}) - f(\rho_s^{g,D})\right] \;>\; 0. \tag{22}$$

Because $\pi_t^{cp} > 0$ at every date, the two regimes never coincide exactly; the statement is asymptotic, not one about a finite transition date. In case (b), and only there, a purely transitory episode of monetary hawkishness leaves a genuinely permanent mark on the long-run carbon stock, even though ϕ_π has no effect on the real allocation once π_t^{cp} has (asymptotically) vanished. For the empirically relevant case $\mu \in (0,1)$ the gap does *not* converge to a positive constant; I refer to this case as a *long-lived*, or persistent, climate legacy, and reserve 'permanent' for the $\mu = 0$ limit. Section 4.5 reports how slowly the gap decays under the baseline calibration, where $\mu = 0.05$.

Remark 3 (Contemporaneous-stance variant)**.** If the portfolio rule were instead driven by the contemporaneous stance indicator r_t of equation (3)–the specification of the previous version of this paper–then $r_t - r^n = (\phi_\pi - 1)\pi_t^{cp}/\Delta(\phi_\pi)$ has slope $(1+\psi\phi_y)\pi_t^{cp}/\Delta(\phi_\pi)^2 > 0$ by algebra alone, so the ordering in Proposition 4 would hold without condition (20), with equality of the two green-capital shares at $t = 0$ and of the lagged carbon stocks at $t = 0$ and $t = 1$ (strict from $t = 2$). Under the timing of Section 3.2 the return relevant to the young is r_{t+1}^{old}, and the monotonicity becomes calibration-dependent, as (20) shows; its sign is nonetheless preserved at the baseline. The two versions coincide in the long-run limit: since $\pi_0^{cp} = 0$, $\sum_{t\geq 0}(r_{t+1}^{old} - r^n) = \sum_{t\geq 0}(r_t - r^n)$ within each regime, so the $\mu = 0$ legacy of Corollary 2 is identical under the two specifications (0.1031 at the baseline); at $t = 30$ the gap is 0.0360 with r_{t+1}^{old} and 0.0394 with r_t.

The proofs are in Appendix D. This is the paper's most policy-relevant result, and it is conditional, as stated, on a policy-invariant cost-push path and on the reduced-form mapping of Assumption 2. It does not say that hawkish monetary policy is a bad idea; it says that its cost is not confined to the conventional output-gap loss captured by $\mathcal{W}(\cdot)$ in Proposition 3–it also includes a climate cost that operates through the timing of green investment, is invisible to a central bank that looks only at $(\pi_t, \hat{y}_t)$, and is long-lived because the carbon stock is a slow-moving, accumulating state variable. Section 6.1 draws out the implication for the joint design of ϕ_π and ξ.

4.5 Numerical Illustration

Figures 2–6 illustrate Propositions 1–4 under the calibration reported in Appendix C, which also reports every numerical value quoted in this section. Figure 2 plots the imported cost-push process of Assumption 1. Figures 3 and 4 plot realised CPI inflation and the output gap along the transition under a dovish ($\phi_\pi = 1.05$), a benchmark ($\phi_\pi = 1.50$), and a hawkish ($\phi_\pi = 4.50$) monetary rule: the more hawkish rule cuts peak realised inflation by about 70.6 percent relative to the benchmark rule (from 3.000 to 0.882 percent at $t = \hat{T}$; Table 4), at the cost of an output gap that is more than twice as large in magnitude at its trough. For contrast, the genuinely permanent ($\mu = 0$) carbon-stock legacy that Corollary 2 reserves for the non-decaying limit would be about 0.10 under this calibration–roughly 2.9 times the finite-horizon gap of 0.036 actually reported in Table 5–which is a further sense in which the calibrated, slowly-decaying case is a materially different, and much less dramatic, object than the limiting case the word 'permanent' describes. Figure 5 plots the cohort losses $\mathcal{L}_1^O(\phi_\pi)$ and $\mathcal{L}_1^Y(\phi_\pi)$ of Proposition 2, evaluated at the point of impact $t^* = 1$ (Section 3.2), against ϕ_π, and marks the closed-form crossing point ϕ_π^* of Corollary 1; under the baseline calibration $\phi_\pi^* = 4.00$, so the hawkish rule in Figures 3–4 sits just beyond the point at which the impact-period incidence of green inflation flips from the old to the young. As Remark 2 makes clear, this crossing point is specific to $t = 1$: at later dates within the rising phase of the hump it falls rapidly toward 1, and beyond the peak $\hat{T}$ it disappears altogether, with the young bearing the larger loss at every hawkish $\phi_\pi > 1$. Figure 6 plots the green-capital-share and carbon-stock paths of Proposition 4 under the dovish and hawkish rules, showing that the hawkish share lies below the dovish share at every date (panel a). Both shares start slightly above the baseline $\bar{\rho} = 0.5$, because at $t = 0$ the return the young collect is depressed below r^n (equation (18)), and both then dip below it as the cost-push hump matures; the dip is earlier and deeper under the hawkish rule (minimum 0.4917 at $t = 5$, against 0.4981 at $t = 9$ under the dovish rule). Because the resulting wedge between the two carbon-stock levels is small relative to the level itself (panel b), panel (c) plots the hawkish-minus-dovish gap directly, making visible the long-lived, only slowly decaying legacy of Corollary 2 at the baseline $\mu = 0.05$: the gap peaks at 0.058 around $t = 14$, is 0.036 at $t = 30$, and is still 0.008 at $t = 60$ and 0.001 at $t = 100$. Recall that one period is one year (Section 3.1).

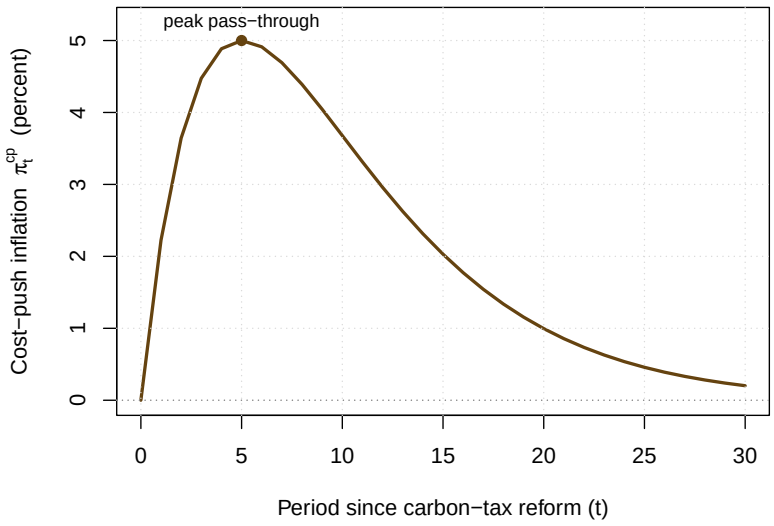


Figure 2. The imported cost-push inflation process π_t^{cp} (Assumption 1)
Source: Author's calculations (*simulation.R*).

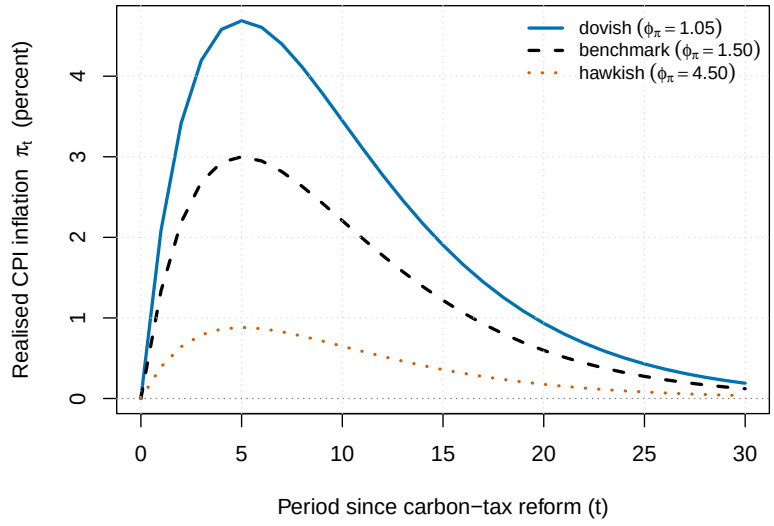


Figure 3. Realised CPI inflation under three monetary regimes (Proposition 1)
Source: Author's calculations (*simulation.R*).

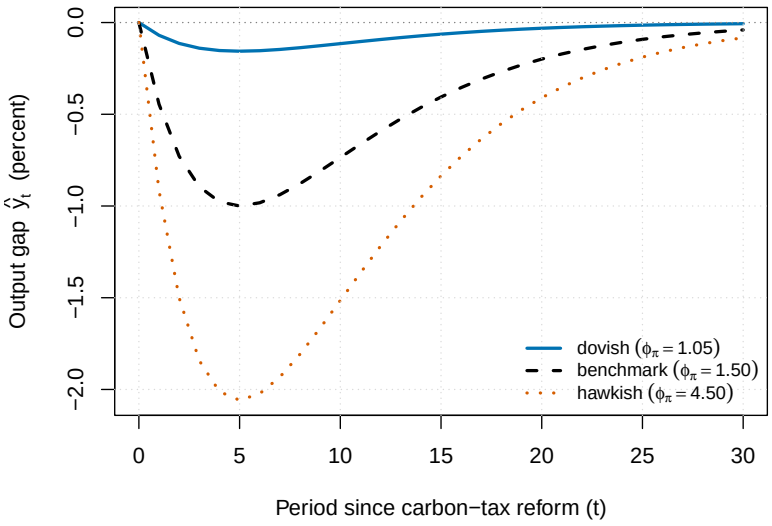


Figure 4. The output gap under three monetary regimes (Proposition 1)
Source: Author's calculations (*simulation.R*).

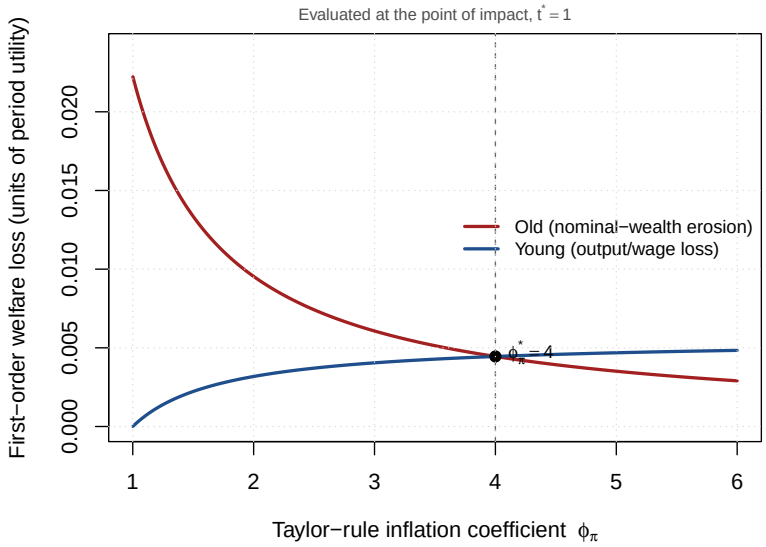


Figure 5. Cohort welfare losses and the monetary-neutral coefficient ϕ_π^* (Proposition 2 and Corollary 1)
Source: Author's calculations (*simulation.R*).

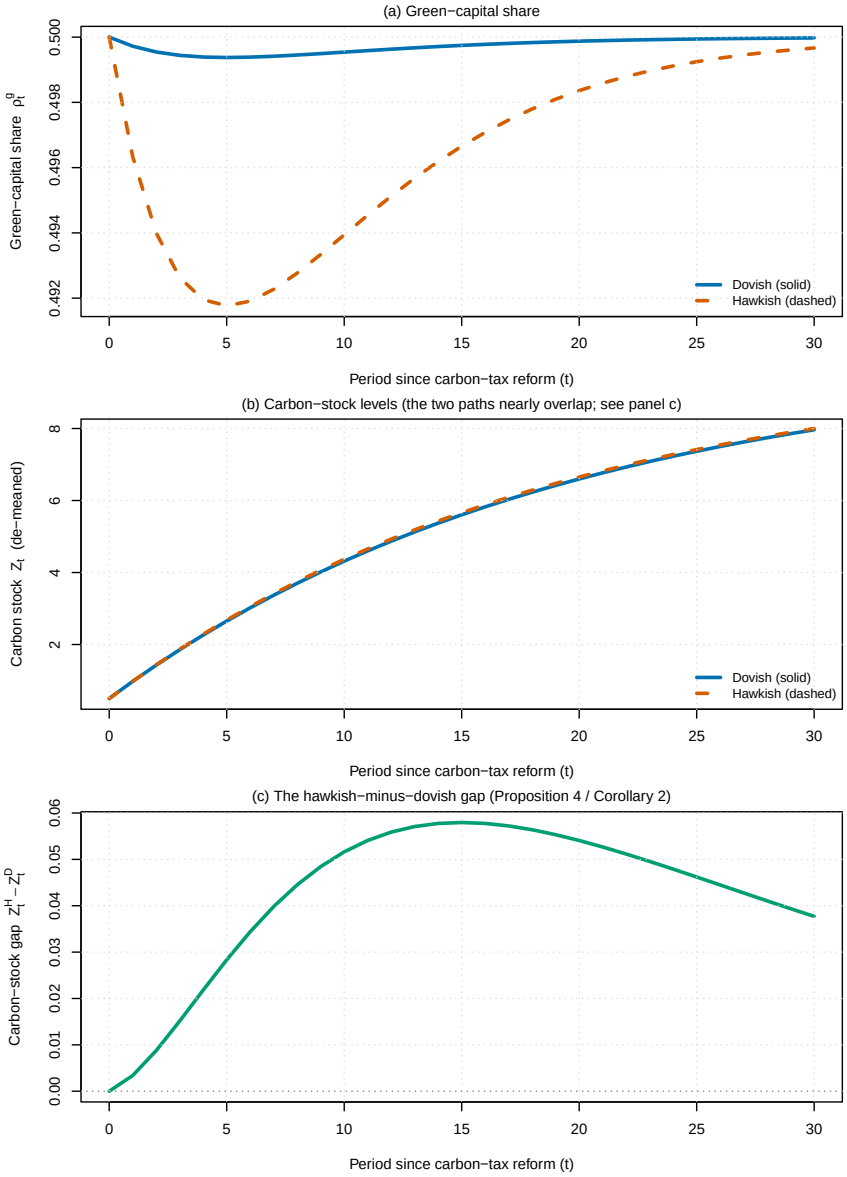


Figure 6. The green-capital share, the carbon stock and the hawkish-minus-dovish carbon-stock gap (Proposition 4 and Corollary 2)
Source: Author's calculations (*simulation.R*). *Note.* Panel (a): green-capital share; panel (b): carbon-stock level; panel (c): hawkish-minus-dovish carbon-stock gap, under dovish versus hawkish monetary policy. Dovish is plotted as a solid blue line and hawkish as a dashed vermillion line throughout; because the gap in panel (b) is small relative to the level of the carbon stock, panel (c) plots it directly.

5. Robustness

5.1 Forward-Looking Price Setting

Equation (5) is static. A natural concern is whether the baseline results survive a forward-looking, New-Keynesian-style Phillips curve,

$$\pi_t \;=\; \kappa \hat{y}_t + \beta\, \pi_{t+1} + \pi_t^{cp}, \qquad \beta \in (0,1). \tag{23}$$

Substituting the Taylor rule (4) into the Fisher relation (3) and then into the IS relation (6) gives, exactly as in the proof of Proposition 1, $\hat{y}_t(1 + \psi\phi_y) = \psi(1 - \phi_\pi)\pi_t$; this step uses only equations (3)–(6) and is unaffected by making equation (5) forward-looking. Substituting this expression for $\hat{y}_t$ into the forward-looking Phillips curve gives

$$\begin{aligned} \pi_t \left[1 - \frac{\kappa\psi(1-\phi_\pi)}{1+\psi\phi_y}\right] \;&=\; \beta\pi_{t+1} + \pi_t^{cp} \quad \Longleftrightarrow \\ \pi_t \;=\; \lambda\, \pi_{t+1} \;+\; \frac{1+\psi\phi_y}{\Delta(\phi_\pi)}\, \pi_t^{cp}, \qquad \lambda \;&\equiv\; \frac{\beta(1+\psi\phi_y)}{\Delta(\phi_\pi)}, \end{aligned} \tag{24}$$

using the same $\Delta(\phi_\pi)$ of equation (11). Because $\{\pi_t^{cp}\}$ is a known, bounded, eventually-decaying deterministic sequence, the unique bounded (non-explosive) solution is the forward sum

$$\pi_t \;=\; \frac{1+\psi\phi_y}{\Delta(\phi_\pi)} \sum_{j=0}^{\infty} \lambda^j\, \pi_{t+j}^{cp}, \qquad \text{provided } |\lambda| < 1, \tag{25}$$

which nests equation (12) exactly at $\beta = 0$, where $\lambda = 0$ and only the contemporaneous term survives. Boundedness requires $|\lambda| < 1$ together with $\Delta(\phi_\pi) > 0$ for sign preservation (as in Proposition 1); together these give

$$\phi_\pi \;>\; 1 - \frac{(1-\beta)(1+\psi\phi_y)}{\kappa\psi}, \tag{26}$$

which nests condition (13) exactly at $\beta = 0$ and is *increasing* in β: the required hawkishness rises toward the classical unit threshold as the Phillips curve becomes more forward-looking, reflecting the additional destabilising channel through which current inflation now also feeds off expected future inflation. This corrects an earlier, unqualified claim that the threshold is 'smaller for $\beta > 0$'; the correct comparative static is that $\phi_\pi^{\min}(\beta)$ rises with β and reaches the classical threshold 1 only in the fully forward-looking limit $\beta \to 1$. This is a boundedness (stability) condition for the forward-looking inflation recursion in the present reduced-form extension. It is not a full rational-expectations determinacy result in the sense of Woodford (2003) or Galí (2015): the IS relation remains static and expectations enter only through the Phillips curve, so I do not claim to have ruled out multiplicity in a canonical forward-looking New Keynesian system. What the extension does show is that the same well-posedness scalar $\Delta(\phi_\pi)$ governs stability once inflation expectations enter, and that the

required hawkishness rises toward the unit threshold as $\beta \to 1$. The sign and monotonicity results of Propositions 1–2 continue to hold qualitatively under this corrected threshold, since they depend only on $\Delta(\phi_\pi)$ entering the same places. I do not re-derive the closed-form crossing point (15) under discounting, both because it requires specifying the entire anticipated path of π_t^{cp} rather than a single value, and because the old's realised bond return of equation (2) would itself need to be re-derived using the forward-looking solution for π_{t-1}; I leave this integration for future work; the same applies to Proposition 4, whose return (18) would then involve the forward sum for π_{t+1}.

5.2 Interest-Rate Smoothing

Replacing equation (4) with a partial-adjustment rule, $i_t = \rho_i\, i_{t-1} + (1-\rho_i)\left[r^n + \phi_\pi \pi_t + \phi_y \hat{y}_t\right]$ for $\rho_i \in [0, 1)$, spreads the same cumulative policy response over more periods. Because i_{t-1} now enters the current-period system directly, the closed form of Proposition 1 no longer applies unchanged: the period-by-period system yields a first-order recursion in $x_t \equiv i_t - r^n$ with autoregressive root

$$c = \frac{\rho_i}{1-(1-\rho_i)\,G}, \qquad G \equiv \frac{\kappa\psi\phi_\pi + \psi\phi_y}{\kappa\psi - 1} \quad (\kappa\psi \neq 1). \tag{27}$$

If $\kappa\psi < 1$, then $G < 0$ and $c \in (0, \rho_i)$: the recursion is bounded for every rule. If $\kappa\psi > 1$, as in the baseline, boundedness of the solution ($|c| < 1$) requires

$$\phi_\pi > \phi_\pi^{\min} + \frac{2\rho_i(\kappa\psi - 1)}{(1-\rho_i)\,\kappa\psi}, \tag{28}$$

so smoothing raises the hawkishness needed, and does so more the more persistent the rule. At the baseline this threshold equals 0.58, 0.92, 1.25 and 4.25 for $\rho_i = 0.25, 0.40, 0.50, 0.80$; the dovish rule $\phi_\pi = 1.05$ is therefore admissible only for $\rho_i < 4/9$. An earlier version of this section stated that smoothing with $\rho_i = 0.5$ leaves every result unchanged; that statement was not supported by the accompanying script and is corrected here: it holds for every rule that satisfies the bound, but not for the baseline dovish rule at $\rho_i = 0.5$, whose recursion is explosive ($c = -1.67$).

Table 8 in Appendix C reports the dovish-versus-hawkish comparison at $\rho_i \in \{0, 0.25, 0.40\}$ at the baseline $\kappa = 2$, and at $\rho_i \in \{0, 0.5, 0.8\}$ with $\kappa = 0.5$ ($\kappa\psi < 1$), where the recursion is bounded for all rules. In every row the signs of the incidence and climate results are unchanged: the hawkish green-capital share lies below the dovish share at every date, and the carbon-stock gap at $t = 30$ is positive and close to its unsmoothed value. Because $i_0 = r^n$, the first-period system coincides with the baseline one with response coefficients scaled by $(1-\rho_i)$, so the crossing point at the point of impact becomes $\phi_\pi^*(\rho_i) = \left[1 + \bar{b}\,(1 + \psi(1-\rho_i)\phi_y)/(\bar{w}\nu\psi)\right]/(1-\rho_i)$, which the simulation reproduces exactly.

5.3 Vanishing Portfolio Sensitivity as a Limiting Case

Proposition 4 operates entirely through the portfolio sensitivity ζ in equation (7). As $\zeta \to 0$ the green-capital share stops responding to the return differential, so the monetary-policy-induced differential has no effect on the allocation: $\rho_t^{g,H} = \rho_t^{g,D} = \bar{\rho}$ for all t, and the legacy gap of Corollary 2 collapses to zero. This is a useful boundary case: it shows that the climate cost of monetary hawkishness is not a generic feature of any model with capital, but specifically a consequence of the responsiveness of portfolio allocation to the return spread between nominal and real assets. It should not be confused with perfect substitutability, which would correspond to an infinitely elastic allocation ($\zeta \to \infty$ in the adjustment-cost interpretation of Section 3.5) and, in a bounded-share version of the rule, to corner allocations; equation (7) is silent about that case, and the interiority condition (21) shows that a sufficiently large ζ would breach the bounds. Within the interior region, the larger ζ, the more forcefully monetary policy spills over into the pace of decarbonisation, as the linear scaling in Table 7 (Panel B) shows.

5.4 Sensitivity of the Monetary-Neutral Coefficient and the Climate-Legacy Gap

Corollary 1 states that ϕ_π^* does not depend on κ or on $\bar{\pi}^{cp}$. Table 6 in Appendix C confirms this numerically by recomputing ϕ_π^* over an extended grid of (κ, ν) combinations that now includes the baseline value $\kappa = 2.00$ (Section 5.5 discusses why this baseline is steep even at annual frequency): the crossing point is exactly invariant to κ and moves only with the wage-output elasticity ν, exactly as equation (15) predicts. This exercise is better described as a numerical verification of an algebraic property of the closed form than as a sensitivity analysis in the usual sense, since invariance to κ is already implied by equation (15) rather than being an independent numerical finding.

A genuine sensitivity exercise instead varies the parameters that actually govern the paper's two channels: $\bar{b}/\bar{w}$, ν, ψ, and ϕ_y for the incidence coefficient ϕ_π^* of equation (15), and ζ and μ for the carbon-stock gap of Corollary 2. Table 7 in Appendix C reports both, expressed as ratios to their own baseline values so as to avoid attaching a false sense of empirical precision to the model's normalised units (Section 5.5): ϕ_π^* is decreasing in ψ and in $\bar{b}/\bar{w}$–a larger interest-sensitivity of demand, or a smaller old-to-young wealth ratio, means less monetary hawkishness is needed to equalise the two cohorts' losses at the point of impact–and increasing in ϕ_y and ν; the carbon-stock gap at $t = 30$ scales linearly in ζ, as equation (7) implies directly, and is decreasing in μ, since a faster-decaying carbon stock erodes any given period's emissions differential more quickly. Panel C varies the Phillips-curve slope: the gap at $t = 30$ is decreasing in κ (ratios 4.09, 2.93, 1.84, 1.00, 0.47 for $\kappa = 0.25,\ 0.5,\ 1,\ 2,\ 4$), because the monetary stance has less leverage over the real return when prices are more flexible; the sign condition (20) and the interiority condition (21) hold for every κ on the grid (Table 9), and the baseline $\kappa = 2$ is thus conservative for the size of the climate legacy.

5.5 Discussion and Limitations

Several simplifications are worth flagging explicitly, in the interest of transparent reporting. First, the aggregate-supply and aggregate-demand relations (5)–(6) are reduced-form rather than derived from optimising price-setters and a consumption Euler equation; this buys tractability within the two-period OLG environment at the cost of the micro-foundations available in a fully specified New Keynesian model. Second, the cost-push process of Assumption 1 is imported as a parametric approximation to the structural model's inflation path rather than re-derived jointly with the monetary block; a fully integrated model would let realised inflation feed back into the real allocation (for instance, through the damage function), which I abstract from here, and Proposition 4 in particular is conditional on this path being unaffected by the monetary stance, as noted in Section 4.4. Proposition 4 is also conditional on the reduced-form mapping $L_t^b = f(\rho_t^g)$ of Assumption 2, which maps a savings share onto an employment share with no production structure behind it; a two-sector block deriving brown-sector activity from optimising firms is beyond the scope of this paper. Third, the portfolio rule of equation (7) is reduced-form, and its sensitivity ζ is a parameter rather than a micro-founded financial friction; Section 5.3 shows the mechanism it drives vanishes as $\zeta \to 0$, and its quantitative magnitude is not disciplined by independent evidence in this paper. Fourth, the model has no aggregate uncertainty, no zero lower bound, and no within-cohort heterogeneity in bond holdings or labour income, all of which are likely to matter for the quantitative–if not the qualitative–size of the effects I identify. Fifth, the cohort welfare losses of Section 3.6 are a within-period comparison between the two cohorts alive at a given date, not a lifetime welfare accounting for a single cohort across its two periods of life; folding the young's own portfolio return into a lifetime criterion would require carrying their asset position forward into next period's $\mathcal{L}_{t+1}^O$, which I leave for future work, though I would not expect it to overturn the direction of the crossing-point result, since it would attenuate rather than reverse the asymmetry between a cohort's exposure to nominal-wealth erosion in old age and to labour-market risk when young. Sixth, the baseline Phillips-curve slope $\kappa = 2.00$ is steep. One model period is one year (Section 3.1), so κ is directly comparable with annual-frequency slopes, and 2.00 lies above what estimates typically support; it also implies $\kappa\psi > 1$, which is why interest-rate smoothing requires the bound of Section 5.2. I keep it as the baseline because Corollary 1 and Table 6 confirm that none of the closed-form incidence results depends on its value, and because Panel C of Table 7 shows that a lower, more conventional κ makes the climate legacy of Proposition 4 larger, not smaller, so the baseline is conservative for that result. Seventh, because $\bar{b}$ and $\bar{w}$ are normalised rather than calibrated to data, the carbon-stock gap and the cohort losses reported in Appendix C are in the model's own de-meaned units and carry no direct empirical scale; Table 7 reports sensitivities as ratios for this reason. I view the model as a deliberately stylised device for isolating the assignment and transition-speed logic of Propositions 3–4, not as a quantitative policy tool in its current form.

6. Conclusions

This paper has added a monetary side to a two-sector overlapping-generations model of the energy transition, asking how the aggressiveness of the central bank's response to carbon-tax-induced green inflation shapes both the intergenerational incidence of the transition and its speed. Four results follow from a common, closed-form solution of a reduced-form Phillips curve, IS relation, and Taylor rule imposed on an imported cost-push process. The condition for the period-by-period system to be well-posed–the augmented green Taylor condition–is generally weaker than the classical unit threshold, because output-gap feedback in the policy rule contributes to stabilisation alongside the direct inflation-feedback channel; Section 5.1 shows that a boundedness (stability) condition for the forward-looking inflation recursion takes the same form once the Phillips curve is made forward-looking, without claiming full rational-expectations determinacy for this reduced-form block. The Taylor-rule coefficient governs the split of the burden of green inflation between old bondholders, who are protected by a hawkish stance, and young wage-earners, who are not; the coefficient that equalises the two losses has a simple closed form, independent of price stickiness and of the size of the shock. Because a lump-sum fiscal transfer can achieve the same equalisation without an output cost, provided sufficient transfer revenue is available, the efficient policy design assigns monetary policy to price and output stabilisation and fiscal transfers to intergenerational equity, rather than leaning on monetary policy for both. Finally, and most importantly for policy, conditional on a policy-invariant cost-push path and on a reduced-form mapping from portfolios to brown-sector activity, a more hawkish stance raises the return the young collect on nominal assets, slows the reallocation of savings toward green capital and leaves a lasting, only slowly decaying mark on the carbon stock, so that the cost of fighting green inflation is not confined to the conventional output-gap loss: it includes a climate cost that is invisible to a central bank that targets $(\pi_t, \hat{y}_t)$ alone.

6.1 Policy Implications

Two implications follow directly from Propositions 3 and 4. First, central banks facing carbon-tax-induced green inflation have a stronger case for accommodation–looking through the cost-push component rather than fighting it with the full force of a conventional Taylor rule–than they would facing a generic demand shock of the same size, precisely because part of the cost of hawkishness here is a long-lived climate cost that a standard stabilisation-only loss function does not price in (conditional, as Proposition 4 states, on a policy-invariant cost-push path and on the reduced-form mapping of Assumption 2). This is not a claim that central banks should ignore inflation; it is a claim that the efficient response to a recognisably transitory, policy-engineered, supply-side inflation process differs from the efficient response to a persistent, demand-driven one, echoing–though in a very different setting–the broader debate on how monetary policy ought to treat inflation that originates outside the central bank's control. Second, distributional concerns about green inflation are better addressed through targeted fiscal transfers, of the kind already studied

in the structural model, than through the calibration of the interest-rate rule itself: asking monetary policy to do the equity job as well as the stabilisation job is feasible, by Proposition 2, but weakly–and, away from the knife-edge case discussed there, strictly–dominated, by Proposition 3, at least as long as the fiscal transfer remains close to lump-sum and sufficient revenue is available. The incidence result should also not be read as a calibration target for the rule: the monetary-neutral coefficient $\phi_\pi^* = 4.00$ is a point-of-impact object, which by Remark 2 falls to about 1.02 within five periods and then ceases to exist once the cost-push process turns down, so it cannot guide the calibration of an interest-rate rule meant to operate over the whole transition–a further reason to leave distributional targeting to fiscal transfers.

Several extensions are natural. A fully dynamic, stochastic version of the model, solved with a proper rational-expectations forward-looking Phillips curve rather than the static reduction used here, would let the cost-push process itself respond to policy rather than being imported as a primitive, closing the loop between the real and nominal blocks. Introducing heterogeneity within cohorts–particularly in bond holdings, following the logic of the bank-lending channel through which monetary policy is transmitted unevenly across borrowers–would let the model speak to within-generation as well as across-generation distributional effects of green monetary policy. A third extension, flagged in Section 4.3, is to re-derive the policy-assignment result of Proposition 3 under a distortionary, rather than lump-sum, fiscal transfer, turning the clean assignment principle into a genuine trade-off disciplined by the relative efficiency costs of the two instruments. Finally, embedding the reduced-form portfolio rule of Section 3.5 in a proper financial-intermediation block, in the spirit of the bank-lending-channel literature, would discipline the parameter ζ that this paper's climate-transition-speed result depends on, and would let the model be taken more seriously as a quantitative, rather than purely qualitative, guide to policy design. I leave these extensions for future work.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, I used Grammarly to improve language and readability, Claude Sonnet (5.0), and DeepSeek (V4) to assist with both language refinement and coding tasks. After using these tools, I reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Annicchiarico, B., & Di Dio, F. (2015). Environmental policy and macroeconomic dynamics in a New Keynesian model. *Journal of Environmental Economics and Management*, *69*, 1–21.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, *109*(6), 2333–2367.

Barrage, L. (2020). Optimal dynamic carbon taxes in a climate-economy model with distortionary fiscal policy. *Review of Economic Studies*, *87*(1), 1–39.

Bondesan, M., Cafferata, A., & Blampied, N. (2026). *The heterogeneous impact of the energy transition* [Available at SSRN 6995838].

Bovenberg, A. L., & Heijdra, B. J. (1998). Environmental tax policy and intergenerational distribution. *Journal of Public Economics*, *67*(1), 1–24.

Clarida, R., Galí, J., & Gertler, M. (1999). The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature*, *37*(4), 1661–1707.

Clarida, R., Galí, J., & Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, *115*(1), 147–180.

Diluiso, F., Annicchiarico, B., Kalkuhl, M., & Minx, J. C. (2021). Climate actions and stranded assets: The role of financial regulation and monetary policy. *Journal of Environmental Economics and Management*, *110*, 102548.

Doepke, M., & Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, *114*(6), 1069–1097.

European Commission, Directorate-General for Economic and Financial Affairs. (2025). *Trends in carbon intensity and the macroeconomic role of the EU emissions trading system* (tech. rep.). European Commission. https://economy-finance.ec.europa.eu/trends-carbon-intensity-and-macroeconomic-role-eu-emissions-trading-system_en

European Environment Agency. (2025). *Use of auctioning revenues generated under the EU emissions trading system* (Indicator assessment). European Environment Agency. https://www.eea.europa.eu/en/analysis/indicators/use-of-auctioning-revenues-generated

Eurostat. (2026). *Inflation in the euro area* (Statistics Explained). Statistical Office of the European Union. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Inflation_in_the_euro_area

Ferrari, A., & Nispi Landi, V. (2024). Whatever it takes to save the planet? central banks and unconventional green policy. *Macroeconomic Dynamics*, *28*(2), 299–324.

Fischer, C., & Springborn, M. (2011). Emissions targets and the real business cycle: Intensity targets versus caps or taxes. *Journal of Environmental Economics and Management*, *62*(3), 352–366.

Fischer, S. (1977). Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of Political Economy*, *85*(1), 191–205.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: An introduction to the New Keynesian framework and its applications* (2nd). Princeton University Press.

Golosov, M., Hassler, J., Krusell, P., & Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, *82*(1), 41–88.

Heutel, G. (2012). How should environmental policy respond to business cycles? optimal policy under persistent productivity shocks. *Review of Economic Dynamics*, *15*(2), 244–264.

International Energy Agency. (2022). *World energy investment 2022* (tech. rep.). IEA. Paris.

International Energy Agency. (2024). *World energy investment 2024* (tech. rep.). IEA. Paris.

International Energy Agency. (2025). *World energy investment 2025* (tech. rep.). IEA. Paris.

John, A., & Pecchenino, R. (1994). An overlapping generations model of growth and the environment. *Economic Journal*, *104*(427), 1393–1410.

Känzig, D. R. (2023). *The unequal economic consequences of carbon pricing* (NBER Working Paper No. 31221). National Bureau of Economic Research.

Kotlikoff, L. J., Kubler, F., Polbin, A., Sachs, J. D., & Scheidegger, S. (2021). Making carbon taxation a generational win win. *International Economic Review*, *62*(1), 3–46.

McKibbin, W. J., Morris, A. C., Panton, A., & Wilcoxen, P. (2020). Climate change and monetary policy: Issues for policy design and modelling. *Oxford Review of Economic Policy*, *36*(3), 579–603.

Mundell, R. A. (1962). The appropriate use of monetary and fiscal policy for internal and external stability. *IMF Staff Papers*, *9*(1), 70–79.

Olovsson, C., & Vestin, D. (2023). *Greenflation?* (Working Paper Series No. 420). Sveriges Riksbank.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, *39*, 195–214.

Tinbergen, J. (1952). *On the theory of economic policy*. North-Holland.

Woodford, M. (2003). *Interest and prices: Foundations of a theory of monetary policy*. Princeton University Press.

Appendix A. Parameter Summary

Table 1. Model parameters and notation (one model period = one year)

Parameter	Description	Restriction
Real block (imported from the BCB model)		
τ_t	Carbon tax at t	$\tau_t = \bar{\tau}(1 - e^{-\partial t})$
π_t^{cp}	Cost-push (green-inflation) process	Assumption 1
$\bar{\pi}^{cp}$	Peak magnitude of π_t^{cp}	$\bar{\pi}^{cp} > 0$
$\hat{T}$	Period at which π_t^{cp} peaks	$\hat{T} > 0$
Monetary block		
ϕ_π	Taylor-rule inflation coefficient	$\phi_\pi \geq 0$
ϕ_y	Taylor-rule output-gap coefficient	$\phi_y \geq 0$
π^*	Inflation target	normalised to 0
r^n	Natural (flexible-price) real rate	$r^n > 0$
κ	Phillips-curve slope	$\kappa > 0$
ψ	IS-curve interest sensitivity	$\psi > 0$
ρ_i	Interest-rate smoothing (Section 5.2)	$\rho_i \in [0, 1)$
β	Forward-looking discount factor (Section 5.1)	$\beta \in (0, 1)$
Cohort incidence		
$\bar{b}$	Old's per-capita real bond wealth	$\bar{b} > 0$
$\bar{w}$	Flexible-price steady-state real wage	$\bar{w} > 0$
ν	Wage-output elasticity	$\nu > 0$
ξ	Transfer share to the old	$\xi \in [0, 1]$
ϕ_π^*	Monetary-neutral coefficient	Corollary 1
Portfolio choice and climate		
ζ	Portfolio sensitivity to the return spread	$\zeta \geq 0$
$\bar{\rho}$	Baseline (frictionless) green-capital share	$\bar{\rho} \in (0, 1)$
ρ_t^g	Equilibrium green-capital share	$\rho_t^g \in (0, 1)$, condition (21)
$f(\cdot)$	Brown-employment mapping, $L_t^b = f(\rho_t^g)$	strictly decreasing, Lipschitz; baseline $1 - \rho$
ε	Emission intensity of brown activity	$\varepsilon > 0$
μ	Atmospheric carbon-stock decay rate	$\mu \in (0, 1)$
Z_t	Carbon stock (de-meaned)	$Z_t \geq 0$
Z_0	Initial carbon stock, common to all regimes	$Z_0 = 0$ in the numerics

Source: Author.

Appendix B. Summary of Key Equations

Table 2. Key equations of the model

Equation	Description	Ref.
$\pi_t^{cp} = \bar{\pi}^{cp}(t/\hat{T})\exp(1 - t/\hat{T})$	Imported cost-push process	(1)
$r_t = i_t - \pi_t$	Contemporaneous policy real rate	(3)
$r_t^{old} = i_{t-1} - \pi_t$	Old's realised bond return	(2)
$i_t = r^n + \phi_\pi \pi_t + \phi_y \hat{y}_t$	Taylor rule	(4)
$\pi_t = \kappa \hat{y}_t + \pi_t^{cp}$	Green Phillips curve	(5)
$\hat{y}_t = -\psi(r_t - r^n)$	IS relation	(6)
$\rho_t^g = \bar{\rho} - \zeta(r_{t+1}^{old} - \mathrm{MPK}_t^g)$	Portfolio choice	(7)
$L_t^b = f(\rho_t^g)$, baseline $f(\rho) = 1 - \rho$	Reduced-form environmental mapping	Assumption 2
$Z_t = (1 - \mu)Z_{t-1} + \varepsilon L_{t-1}^b$	Carbon-stock accumulation	(8)
$\mathcal{L}_t^O = \bar{b}\,(\pi_t - \phi_\pi \pi_{t-1} - \phi_y \hat{y}_{t-1}) - \xi\, T_t/P_t$	Old's welfare loss	(9)
$\mathcal{L}_t^Y = -\bar{w}\nu\hat{y}_t - (1 - \xi)T_t/P_t$	Young's welfare loss	(10)
$\Delta(\phi_\pi) = 1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1)$	Well-posedness scalar	(11)
$\pi_t = (1 + \psi\phi_y)\pi_t^{cp}/\Delta(\phi_\pi)$	Equilibrium inflation	(12)
$\hat{y}_t = \psi(1 - \phi_\pi)\pi_t^{cp}/\Delta(\phi_\pi)$	Equilibrium output gap	(12)
$\phi_\pi^{\min} = 1 - (1 + \psi\phi_y)/(\kappa\psi)$	Augmented Taylor condition	(13)
$\phi_\pi^* = 1 + \bar{b}(1 + \psi\phi_y)/(\bar{w}\nu\psi)$	Monetary-neutral coefficient (at $t^* = 1$)	(15)
$r_{t+1}^{old} - r^n = N_t/\Delta(\phi_\pi)$	Return the young collect	(18)
$N_t = (\phi_\pi + \psi\phi_y)\pi_t^{cp} - (1 + \psi\phi_y)\pi_{t+1}^{cp}$	Numerator of the return gap	(18)
$\kappa\psi \le 1/(1 - e^{-1/\hat{T}})$	Sign condition for Proposition 4	(20)
$\zeta\bar{\pi}^{cp}[\max\{\phi_\pi, 1\} + \psi\phi_y]/\Delta(\phi_\pi) < \min\{\bar{\rho}, 1 - \bar{\rho}\}$	Interiority (sufficient)	(21)

Source: Author.

Appendix C. Calibration and Numerical Results

Table 3 reports the baseline calibration used to produce Figures 2–6 and every numerical value quoted in Section 4. One model period is one year (Section 3.1). All figures and numbers are produced by the accompanying R script, *simulation.R*, which writes the full numerical log to *numerical_results.txt*; *motivation_figure.R* produces Figure 1. For reproducibility I report the following implementation details.

Initial condition. The pre-reform carbon stock is $Z_0 = 0$ in the de-meaned units of the model, common to all regimes.

Timing convention. The carbon stock follows the lagged law of motion (8), $Z_t = (1-\mu)Z_{t-1} + \varepsilon L^b_{t-1}$, in the code, the proofs and the tables alike. The previous version of the script accumulated emissions contemporaneously, which was inconsistent with equation (8) and Proposition 4; this is corrected, and Tables 5 and 7 are recomputed.

Portfolio rule. The green-capital share is driven by $r^{old}_{t+1} = i_t - \pi_{t+1}$ (equation (7)), and the code no longer clips it to [0, 1]: Table 9 verifies the interiority condition (21) and the sign condition (20) in the baseline and in every sensitivity exercise.

Numerical horizon. Figures and tables run over $t = 0, \ldots, 30$ years. The infinite sum in Corollary 2(b) is approximated by truncating at 2,000 periods, which is exact at the reported precision because π^{cp}_t decays geometrically; the cost-push process itself is never truncated.

Evaluation points. The losses at $\phi_\pi = 1$ and $\phi_\pi = 6$ are evaluated exactly at those values. The previous version evaluated $\mathcal{L}^O$ at $\phi_\pi = 1.001$, which explains the earlier entry 0.02223 in place of the exact 0.02226 implied by equation (1) ($\pi^{cp}_1 = 0.02226$) and the closed form.

The script, together with the full numerical log, accompanies this resubmission as a replication package; a persistent public link will be added at the production stage.

Table 3. Baseline calibration (one model period = one year)

Parameter	Description	Value
ψ	IS-curve interest sensitivity	1.00
κ	Phillips-curve slope	2.00
ϕ_y	Taylor-rule output-gap coefficient	0.50
π^*	Inflation target	0.00
r^n	Natural real rate (per year)	0.02
ν	Wage-output elasticity	0.50
$\bar{b}$	Old's real bond wealth (normalised)	1.00
$\bar{w}$	Steady-state real wage (normalised)	1.00
ζ	Portfolio sensitivity to the return spread	0.40
$\bar{\rho}$	Baseline green-capital share	0.50
ε	Emission intensity	1.00
μ	Carbon-stock decay rate (per year)	0.05
Z_0	Initial carbon stock (de-meaned)	0.00
$\bar{\pi}^{cp}$	Peak cost-push inflation	0.05 (5%)
$\hat{T}$	Peak-incidence period (years)	5

Source: Author.

Table 4. Realised inflation and the output gap by monetary regime, at $t = \hat{T} = 5$

Regime	ϕ_π	Peak π_t (%)	Trough $\hat{y}_t$ (%)
Dovish	1.05	4.688	−0.156
Benchmark	1.50	3.000	−1.000
Hawkish	4.50	0.882	−2.059

Source: Author's calculations (*simulation.R*). *Note.* Existence/well-posedness threshold: $\phi_\pi^{\min} = 0.250$; all three regimes satisfy $\Delta(\phi_\pi) > 0$ (Proposition 1).

Table 5. Numerical results for Propositions 2 and 4

Quantity	Value
Old's loss at $t^* = 1$, $\phi_\pi = 1$ (exact), $\mathcal{L}_1^O(1)$	0.02226
Young's loss at $t^* = 1$, $\phi_\pi = 6$, $\mathcal{L}_1^Y(6)$	0.00484
Monetary-neutral coefficient, ϕ_π^* (numerical, $t^* = 1$)	4.0000
Monetary-neutral coefficient, ϕ_π^* (closed form, eq. (15))	4.0000
Common loss at ϕ_π^*, $t^* = 1$	0.00445
General-t crossing point $\phi_\pi^*(t)$, $t = 2, 3, 4, 5$	1.526, 1.212, 1.089, 1.023
Green-capital share at $t = 0$, dovish / hawkish	0.5083 / 0.5016
Minimum green-capital share, dovish regime (at $t = 9$)	0.4981
Minimum green-capital share, hawkish regime (at $t = 5$)	0.4917
Carbon stock at $t = 30$, dovish regime	7.8609
Carbon stock at $t = 30$, hawkish regime	7.8970
Finite-horizon carbon-stock gap, $Z_{30}^H - Z_{30}^D$	0.0360
Peak of the gap over $t = 0, \ldots, 30$ (at $t = 14$)	0.0584
Gap at $t = 60$ / $t = 100$	0.0081 / 0.0011
$\mu = 0$ limit of the gap, $\varepsilon \sum_{s \geq 0}(L_s^{b,H} - L_s^{b,D})$	0.1031

Source: Author's calculations (*simulation.R*). *Note.* The old's and young's losses are evaluated at the point of impact $t^* = 1$ (Section 4.2), and $\mathcal{L}_1^O(1)$ is evaluated exactly at $\phi_\pi = 1$. Green-capital shares follow the portfolio rule (7) driven by r_{t+1}^{old}, and carbon stocks the lagged law of motion (8) with $Z_0 = 0$; the hawkish/dovish regimes are $\phi_\pi = 4.50$ and 1.05. The reported gap is the finite-horizon value at $t = 30$, not the $t \to \infty$ limit, which is zero at the baseline $\mu = 0.05$ (Corollary 2); the $\mu = 0$ row is the strictly positive limit that Corollary 2(b) reserves for the non-decaying case.

Table 6. Robustness of the monetary-neutral coefficient ϕ_π^* to the Phillips-curve slope and the wage-output elasticity

κ	ν	ϕ_π^*
0.15	0.25	7.000
0.30	0.25	7.000
0.60	0.25	7.000
2.00	0.25	7.000
0.15	0.50	4.000
0.30	0.50	4.000
0.60	0.50	4.000
2.00	0.50	4.000
0.15	1.00	2.500
0.30	1.00	2.500
0.60	1.00	2.500
2.00	1.00	2.500

Source: Author's calculations (*simulation.R*). *Note.* The coefficient ϕ_π^* is numerically invariant to the Phillips-curve slope κ and varies only with the wage-output elasticity ν, exactly as predicted by the closed-form expression in Corollary 1 (Section 5.4). The value $\kappa = 2.00$ is the paper's baseline (Table 3) and is included, in addition to the lower range more typical of quarterly Phillips-curve estimates, to confirm the invariance holds at the calibration actually used (Section 5.5).

Table 7. Sensitivity of the monetary-neutral coefficient and of the carbon-stock gap

Panel A: incidence coefficient ϕ_π^* (eq. (15)), baseline = 4.000	
Parameter change	ϕ_π^*
$\psi \times 0.5$	6.000
$\psi \times 2$	3.000
$\phi_y \times 0.5$	3.500
$\phi_y \times 2$	5.000
$\bar{b}/\bar{w} \times 0.5$	2.500
$\bar{b}/\bar{w} \times 2$	7.000
$\nu \times 0.5$	7.000
$\nu \times 2$	2.500
Panel B: carbon-stock gap $Z_{30}^H - Z_{30}^D$, relative to own baseline	
Parameter change	Ratio to baseline
$\zeta \times 0.5$	0.50×
$\zeta \times 2$	2.00×
$\mu \times 0.5$	1.67×
$\mu \times 2$	0.39×
Panel C: Phillips-curve slope, carbon-stock gap $Z_{30}^H - Z_{30}^D$, relative to own baseline	
Parameter change	Ratio to baseline
$\kappa = 0.25$	4.09×
$\kappa = 0.50$	2.93×
$\kappa = 1.00$	1.84×
$\kappa = 2.00$ (baseline)	1.00×
$\kappa = 4.00$	0.47×

Source: Author's calculations (*simulation.R*). *Note.* Genuine sensitivity (Section 5.4), as opposed to the numerical-verification exercise of Table 6. Panel A varies, one at a time, the parameters entering the closed-form incidence coefficient (15), holding the rest at their Table 3 values; because ϕ_π^* is exact and linear in these ratios, the reported values follow directly from equation (15). Panel B varies the two parameters governing the portfolio-substitution channel of Proposition 4 and reports the resulting finite-horizon carbon-stock gap at $t = 30$ as a ratio to its own baseline value (computed with the forward-return portfolio rule and the lagged carbon stock; Section 5.5 explains why ratios, rather than levels, are reported here), holding $\phi_\pi^H = 4.50$, $\phi_\pi^D = 1.05$ fixed. Panel C varies the Phillips-curve slope κ in the same way; the sign and interiority conditions hold for every κ on the grid (Table 9).

Table 8. Interest-rate smoothing: dovish versus hawkish rule

κ	ρ_i	Root c		ϕ_π^*	Peak π_t	Trough $\hat{y}_t$	$\rho_t^{g,H} < \rho_t^{g,D}$	$Z_{30}^H - Z_{30}^D$
		dovish	hawkish	($t^* = 1$)	hawkish (%)	hawkish (%)	all t	
2.0	0	0	0	4.00	0.882	−2.059	yes	0.0360
2.0	0.25	−0.263	−0.041	5.00	0.887	−2.056	yes	0.0353
2.0	0.40	−0.714	−0.085	6.00	0.897	−2.055	yes	0.0345
0.5	0	0	0	4.00	2.308	−5.385	yes	0.1057
0.5	0.50	0.247	0.133	7.00	2.368	−5.316	yes	0.1049
0.5	0.80	0.567	0.381	16.00	3.059	−5.263	yes	0.1023

Source: Author's calculations (*simulation.R*). *Note.* Section 5.2: dovish ($\phi_\pi = 1.05$) versus hawkish ($\phi_\pi = 4.50$) rule under the partial-adjustment rule with smoothing parameter ρ_i. "Root c" is the autoregressive root of $i_t - r^n$; $|c| < 1$ is required for a bounded solution. At the baseline $\kappa = 2$ ($\kappa\psi > 1$) the dovish rule is not admissible for $\rho_i \geq 4/9$ (e.g. $c = -1.67$ at $\rho_i = 0.5$), so the rows with stronger smoothing use $\kappa = 0.5$ ($\kappa\psi < 1$), where the recursion is bounded for every rule; levels are not comparable across the two κ blocks. ϕ_π^* is the crossing point at the point of impact. Green-capital shares use the forward-return portfolio rule, and the last column uses the lagged carbon stock.

Table 9. Interiority and sign checks

Case	$\zeta \sup_t \lvert r_{t+1}^{old} - r^n \rvert$		Limit	$S_t > 0$	$\rho_t^{g,H} < \rho_t^{g,D}$
	dovish	hawkish	$\min\{\bar{\rho}, 1-\bar{\rho}\}$	all t	all t
Baseline ($\zeta = 0.40$, $\kappa = 2$)	0.0083	0.0083	0.5	yes	yes
$\zeta \times 0.5$	0.0042	0.0041	0.5	yes	yes
$\zeta \times 2$	0.0167	0.0166	0.5	yes	yes
$\kappa = 0.25$	0.0088	0.0297	0.5	yes	yes
$\kappa = 0.50$	0.0088	0.0217	0.5	yes	yes
$\kappa = 1.00$	0.0086	0.0141	0.5	yes	yes
$\kappa = 4.00$	0.0079	0.0046	0.5	yes	yes

Source: Author's calculations (*simulation.R*). *Note.* Section 4.4. The entries are the exact values of $\zeta \sup_t |r_{t+1}^{old} - r^n|$, computed over $t = 0, \ldots, 400$, for the dovish (1.05) and hawkish (4.50) rules; the interior region requires them to lie below $\min\{\bar{\rho}, 1-\bar{\rho}\} = 0.5$. At the baseline the largest admissible ζ is about 24 for both rules. $S_t > 0$ is the slope condition of equation (19), checked over $t = 0, \ldots, 2{,}000$, and the last column confirms the ordering of Proposition 4(iii) over $t = 0, \ldots, 30$. The sign condition (20) holds for all cases ($\kappa\psi \leq 4 < 5.52$).

Appendix D. Mathematical Proofs

D.1 Proof of Proposition 1

Proof. Substitute the Fisher relation (3) into the Taylor rule (4) to obtain

$$r_t \;=\; i_t - \pi_t \;=\; r^n + \phi_\pi \pi_t + \phi_y \hat{y}_t - \pi_t. \tag{29}$$

Substitute this into the IS relation (6):

$$\hat{y}_t \;=\; -\psi\left[\phi_\pi \pi_t + \phi_y \hat{y}_t - \pi_t\right] \;=\; \psi(1-\phi_\pi)\pi_t - \psi\phi_y \hat{y}_t, \tag{30}$$

so that

$$\hat{y}_t(1+\psi\phi_y) \;=\; \psi(1-\phi_\pi)\pi_t. \tag{31}$$

Substitute $\hat{y}_t$ from equation (31) into the Phillips curve (5):

$$\pi_t \;=\; \kappa\hat{y}_t + \pi_t^{cp} \;=\; \frac{\kappa\psi(1-\phi_\pi)}{1+\psi\phi_y}\pi_t + \pi_t^{cp}. \tag{32}$$

Collecting terms in π_t,

$$\pi_t\left[1 - \frac{\kappa\psi(1-\phi_\pi)}{1+\psi\phi_y}\right] \;=\; \pi_t^{cp} \quad\Longleftrightarrow\quad \pi_t \cdot \frac{(1+\psi\phi_y) - \kappa\psi(1-\phi_\pi)}{1+\psi\phi_y} \;=\; \pi_t^{cp}. \tag{33}$$

The numerator of the bracketed coefficient is exactly $\Delta(\phi_\pi) \;=\; 1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1)$ of equation (11), so

$$\pi_t \;=\; \frac{(1+\psi\phi_y)\,\pi_t^{cp}}{\Delta(\phi_\pi)}, \tag{34}$$

which is well defined and finite if and only if $\Delta(\phi_\pi) \neq 0$, establishing existence and uniqueness of π_t (and, through equation (31), of $\hat{y}_t$) as claimed in equation (12). Since $1 + \psi\phi_y > 0$ by $\psi, \phi_y \geq 0$, the sign of π_t matches the sign of π_t^{cp} if and only if $\Delta(\phi_\pi) > 0$, which is condition (13) once rearranged:

$$\Delta(\phi_\pi) > 0 \;\Longleftrightarrow\; 1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1) > 0 \;\Longleftrightarrow\; \phi_\pi > 1 - \frac{1+\psi\phi_y}{\kappa\psi}, \tag{35}$$

using $\kappa\psi > 0$ to divide without reversing the inequality. This proves equation (13) and the well-posedness claim. □

D.2 Proof of Proposition 2

Proof. By Proposition 1 evaluated at $t = t^* = 1$ with $\pi_*^{cp} = \pi_1^{cp} > 0$ and $\pi^* = 0$,

$$\pi_1(\phi_\pi) \;=\; \frac{(1+\psi\phi_y)\pi_*^{cp}}{\Delta(\phi_\pi)}, \qquad \hat{y}_1(\phi_\pi) \;=\; \frac{\psi(1-\phi_\pi)\pi_*^{cp}}{\Delta(\phi_\pi)}. \tag{36}$$

Because $\pi_0 = \hat{y}_0 = 0$ (Section 3.2), equation (9) gives $\mathcal{L}_1^O = \bar{b}\,\pi_1$ exactly, with no approximation, so that $\mathcal{L}^O(\phi_\pi) = \bar{b}(1 + \psi\phi_y)\pi_*^{cp}/\Delta(\phi_\pi)$ and $\mathcal{L}^Y(\phi_\pi) = \bar{w}\nu\psi(\phi_\pi - 1)\pi_*^{cp}/\Delta(\phi_\pi)$; this is the sense in which the choice $t^* = 1$, rather than an arbitrary date, is what makes the one-term expression exact (Section 4.2).

(i) Monotonicity of $\mathcal{L}^O$. Since $\Delta(\phi_\pi) = 1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1)$ is strictly increasing in ϕ_π (its derivative is $\kappa\psi > 0$) and $\mathcal{L}^O(\phi_\pi)$ is a positive constant divided by $\Delta(\phi_\pi)$, $\mathcal{L}^O$ is strictly decreasing in ϕ_π. At $\phi_\pi = 1$, $\Delta(1) = 1 + \psi\phi_y$, so $\mathcal{L}^O(1) = \bar{b}\,\pi_*^{cp} > 0$. As $\phi_\pi \to \infty$, $\Delta(\phi_\pi) \sim \kappa\psi\phi_\pi \to \infty$, so $\mathcal{L}^O(\phi_\pi) \to 0$.

(ii) Monotonicity of $\mathcal{L}^Y$. Write $\mathcal{L}^Y(\phi_\pi) = \bar{w}\nu\psi\pi_*^{cp} \cdot g(\phi_\pi)$ with $g(\phi_\pi) \equiv (\phi_\pi - 1)/\Delta(\phi_\pi)$. Then

$$\begin{aligned} g'(\phi_\pi) &= \frac{\Delta(\phi_\pi) - (\phi_\pi - 1)\Delta'(\phi_\pi)}{\Delta(\phi_\pi)^2} \\ &= \frac{\left[1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1)\right] - (\phi_\pi - 1)\kappa\psi}{\Delta(\phi_\pi)^2} \\ &= \frac{1 + \psi\phi_y}{\Delta(\phi_\pi)^2} > 0, \end{aligned} \tag{37}$$

so g, and hence $\mathcal{L}^Y$, is strictly increasing in ϕ_π wherever $\Delta(\phi_\pi) > 0$. At $\phi_\pi = 1$, $g(1) = 0$, so $\mathcal{L}^Y(1) = 0$. As $\phi_\pi \to \infty$, $g(\phi_\pi) = (\phi_\pi - 1)/[1 + \psi\phi_y + \kappa\psi(\phi_\pi - 1)] \to 1/(\kappa\psi)$, so $\mathcal{L}^Y(\phi_\pi) \to \bar{w}\nu\psi\pi_*^{cp}/(\kappa\psi) = \bar{w}\nu\pi_*^{cp}/\kappa$.

(iii) Existence and uniqueness of the crossing point. Define $h(\phi_\pi) \equiv \mathcal{L}^O(\phi_\pi) - \mathcal{L}^Y(\phi_\pi)$ on $(1, \infty)$. By parts (i)–(ii), h is continuous and strictly decreasing (a decreasing function minus an increasing function), with $h(1^+) = \mathcal{L}^O(1) - 0 > 0$ and $\lim_{\phi_\pi \to \infty} h(\phi_\pi) = 0 - \bar{w}\nu\pi_*^{cp}/\kappa < 0$. By the intermediate value theorem, there exists $\phi_\pi^* \in (1, \infty)$ with $h(\phi_\pi^*) = 0$; strict monotonicity of h makes ϕ_π^* unique. Since h is strictly decreasing, $h(\phi_\pi) > 0$ (the old bear the larger loss) for $\phi_\pi < \phi_\pi^*$ and $h(\phi_\pi) < 0$ (the young bear the larger loss) for $\phi_\pi > \phi_\pi^*$. □

D.3 Proof of Corollary 1

Proof. At $\phi_\pi = \phi_\pi^*$, $\mathcal{L}^O(\phi_\pi^*) = \mathcal{L}^Y(\phi_\pi^*)$ requires

$$\frac{\bar{b}(1 + \psi\phi_y)\pi_*^{cp}}{\Delta(\phi_\pi^*)} = \frac{\bar{w}\nu\psi(\phi_\pi^* - 1)\pi_*^{cp}}{\Delta(\phi_\pi^*)}. \tag{38}$$

Because $\pi_*^{cp} > 0$ and $\Delta(\phi_\pi^*) > 0$ appear as common factors on both sides, they cancel, leaving

$$\bar{b}(1 + \psi\phi_y) = \bar{w}\nu\psi(\phi_\pi^* - 1) \quad \Longleftrightarrow \quad \phi_\pi^* = 1 + \frac{\bar{b}(1 + \psi\phi_y)}{\bar{w}\nu\psi}, \tag{39}$$

which is equation (15). Since κ never appears in this expression (it cancelled together with $\Delta(\phi_\pi^*)$), and π_*^{cp} cancelled as well, ϕ_π^* is independent of both. □

D.4 Proof of Proposition 3

Proof. (i) Existence and uniqueness of $\xi^{**}(\phi_\pi)$. The condition $\tilde{\mathcal{L}}_s^O = \tilde{\mathcal{L}}_s^Y$ reads

$$\mathcal{L}_s^O - \xi\frac{T_t}{P_t} \;=\; \mathcal{L}_s^Y - (1-\xi)\frac{T_t}{P_t}, \tag{40}$$

which rearranges to a single linear equation in ξ:

$$\mathcal{L}_s^O - \mathcal{L}_s^Y \;=\; \xi\frac{T_t}{P_t} - (1-\xi)\frac{T_t}{P_t} \;=\; (2\xi-1)\frac{T_t}{P_t} \quad\Longleftrightarrow\quad \xi^{**}(\phi_\pi) \;=\; \frac{1}{2} + \frac{\mathcal{L}_s^O(\phi_\pi) - \mathcal{L}_s^Y(\phi_\pi)}{2\,T_t/P_t}. \tag{41}$$

This solution exists and is unique for any ϕ_π, since the equation is linear with a nonzero coefficient on ξ (namely $2\,T_t/P_t \neq 0$ for $T_t > 0$). It lies in the interior $(0,1)$ if and only if $|\mathcal{L}_s^O(\phi_\pi) - \mathcal{L}_s^Y(\phi_\pi)| < T_t/P_t$, that is, if and only if the available transfer revenue exceeds the incidence gap; this is the feasibility condition of the proposition.

(ii) Separation of the joint optimisation (whenever the feasibility condition holds). Substituting $\xi = \xi^{**}(\phi_\pi)$ from part (i) makes $\tilde{\mathcal{L}}_s^O(\phi_\pi, \xi^{**}(\phi_\pi)) - \tilde{\mathcal{L}}_s^Y(\phi_\pi, \xi^{**}(\phi_\pi)) = 0$ identically, for every ϕ_π, by construction. Hence the distributional term of $\mathcal{W}(\phi_\pi, \xi)$ can be driven to exactly zero for any ϕ_π by an appropriate choice of ξ, and doing so is optimal for any $\omega_d > 0$ since the term is a squared deviation bounded below by zero. The policymaker's problem therefore decomposes as

$$\begin{aligned} \min_{\phi_\pi,\xi} \mathcal{W}(\phi_\pi, \xi) \;&=\; \min_{\phi_\pi}\left[\tfrac{1}{2}\left(\pi_t(\phi_\pi)^2 + \omega_y\hat{y}_t(\phi_\pi)^2\right) + \min_{\xi}\omega_d\left(\tilde{\mathcal{L}}_s^O - \tilde{\mathcal{L}}_s^Y\right)^2\right] \\ &=\; \min_{\phi_\pi}\tfrac{1}{2}\left(\pi_t(\phi_\pi)^2 + \omega_y\hat{y}_t(\phi_\pi)^2\right), \end{aligned} \tag{42}$$

where the inner minimisation over ξ attains zero at $\xi = \xi^{**}(\phi_\pi)$ regardless of ϕ_π, and therefore drops out of the outer problem entirely: the optimal ϕ_π solves the stabilisation-only problem and does not depend on ω_d. □

D.5 Proof of Proposition 4

Proof. Derivation of (18). By the Taylor rule (4) and equation (12),

$$i_t - r^n = \phi_\pi\pi_t + \phi_y\hat{y}_t = \frac{\phi_\pi(1+\psi\phi_y) + \phi_y\psi(1-\phi_\pi)}{\Delta(\phi_\pi)}\,\pi_t^{cp} = \frac{(\phi_\pi + \psi\phi_y)\,\pi_t^{cp}}{\Delta(\phi_\pi)}, \tag{43}$$

and, applying equation (12) at $t+1$, $\pi_{t+1} = (1+\psi\phi_y)\pi_{t+1}^{cp}/\Delta(\phi_\pi)$. Hence $r_{t+1}^{old} - r^n = i_t - \pi_{t+1} - r^n$ is given by (18).

(i) Slope. Write $a \equiv 1+\psi\phi_y$ and $N_t(\phi_\pi) \equiv (\phi_\pi + \psi\phi_y)\pi_t^{cp} - a\,\pi_{t+1}^{cp}$, so that $\partial N_t/\partial\phi_\pi = \pi_t^{cp}$ and $\partial\Delta/\partial\phi_\pi = \kappa\psi$. Then

$$\frac{\partial}{\partial\phi_\pi}\frac{N_t}{\Delta} = \frac{\pi_t^{cp}\Delta - \kappa\psi N_t}{\Delta^2}, \tag{44}$$

and, using $\Delta = a + \kappa\psi(\phi_\pi - 1)$,

$$\begin{aligned}\pi_t^{cp}\Delta - \kappa\psi N_t &= \pi_t^{cp}\left[a + \kappa\psi(\phi_\pi - 1) - \kappa\psi(\phi_\pi + \psi\phi_y)\right] + \kappa\psi\, a\, \pi_{t+1}^{cp} \\ &= a\left[(1-\kappa\psi)\pi_t^{cp} + \kappa\psi\, \pi_{t+1}^{cp}\right] = a\, S_t,\end{aligned} \tag{45}$$

because $a - \kappa\psi - \kappa\psi \cdot \psi\phi_y = a(1 - \kappa\psi)$. This is (19). At $t = 0$, $\pi_0^{cp} = 0$ so $S_0 = \kappa\psi\, \pi_1^{cp} > 0$. For $t \geq 1$, Assumption 1 gives $q_t \equiv \pi_{t+1}^{cp}/\pi_t^{cp} = \frac{t+1}{t}e^{-1/\hat{T}}$, which is strictly decreasing in t with limit $e^{-1/\hat{T}}$, and $S_t = \pi_t^{cp}[1 - \kappa\psi + \kappa\psi\, q_t]$. If $\kappa\psi(1 - e^{-1/\hat{T}}) \leq 1$, then $S_t = \pi_t^{cp}\{[1 - \kappa\psi(1 - e^{-1/\hat{T}})] + \kappa\psi(q_t - e^{-1/\hat{T}})\} > 0$ for every $t \geq 1$, since $q_t > e^{-1/\hat{T}}$ strictly. If instead $\kappa\psi(1-e^{-1/\hat{T}}) > 1$, then $1 - \kappa\psi + \kappa\psi q_t \to 1 - \kappa\psi(1 - e^{-1/\hat{T}}) < 0$, so $S_t < 0$ for all sufficiently large t. Hence $S_t > 0$ for every $t \geq 0$ if and only if condition (20) holds.

(ii) Ordering of returns. Because Δ is strictly increasing in ϕ_π and $\Delta(\phi_\pi^D) > 0$, $\Delta(\phi_\pi) > 0$ on $[\phi_\pi^D, \phi_\pi^H]$. Since S_t does not depend on ϕ_π,

$$r_{t+1}^{old,H} - r_{t+1}^{old,D} = \int_{\phi_\pi^D}^{\phi_\pi^H} \frac{a\, S_t}{\Delta(\phi)^2}\, d\phi > 0 \qquad \text{for every } t \geq 0 \text{ under (20).} \tag{46}$$

(iii) Green-capital share and brown employment. Evaluating the portfolio rule (7) at $\mathrm{MPK}_t^g = r^n$,

$$\rho_t^{g,H} - \rho_t^{g,D} = -\zeta\left(r_{t+1}^{old,H} - r_{t+1}^{old,D}\right) < 0 \qquad \text{for every } t \geq 0 \text{ if } \zeta > 0, \tag{47}$$

with both shares in $(0, 1)$ by the interiority assumption. Since f is strictly decreasing (Assumption 2), $L_t^{b,H} = f(\rho_t^{g,H}) > f(\rho_t^{g,D}) = L_t^{b,D}$.

(iv) Carbon stock. I show $Z_t^H \geq Z_t^D$ by induction on the law of motion (8). At $t = 0$, $Z_0^H = Z_0^D = Z_0$ by the common initial condition. For $t \geq 1$,

$$Z_t^H - Z_t^D = (1 - \mu)\left(Z_{t-1}^H - Z_{t-1}^D\right) + \varepsilon\left(L_{t-1}^{b,H} - L_{t-1}^{b,D}\right), \tag{48}$$

where $(1 - \mu) \in (0, 1)$, the first term is non-negative by the inductive hypothesis, and the second is strictly positive by part (iii) when $\zeta > 0$. Hence $Z_t^H > Z_t^D$ for every $t \geq 1$; in particular $Z_1^H - Z_1^D = \varepsilon(L_0^{b,H} - L_0^{b,D}) > 0$, because emissions of period 0 already differ across regimes ($S_0 > 0$), and the stock at $t = 0$ has not yet had a chance to reflect any difference. With $\zeta = 0$ all paths coincide.

Interiority bound (21). The numerator of (18) is $A_t - B_t$ with $A_t = (\phi_\pi + \psi\phi_y)\pi_t^{cp} \geq 0$ and $B_t = a\pi_{t+1}^{cp} \geq 0$, so $|A_t - B_t| \leq \max\{A_t, B_t\} \leq [\max\{\phi_\pi, 1\} + \psi\phi_y]\,\bar{\pi}^{cp}$ since $\pi_t^{cp} \leq \bar{\pi}^{cp}$. Therefore $\zeta \sup_t |r_{t+1}^{old} - r^n| \leq \zeta[\max\{\phi_\pi, 1\} + \psi\phi_y]\bar{\pi}^{cp}/\Delta(\phi_\pi)$, and condition (21) implies $|\rho_t^g - \bar{\rho}| < \min\{\bar{\rho}, 1 - \bar{\rho}\}$, i.e. $\rho_t^g \in (0, 1)$, at every date. □

D.6 Proof of Corollary 2

Proof. Let $e_t \equiv \varepsilon[f(\rho_t^{g,H}) - f(\rho_t^{g,D})]$, which is strictly positive for every $t \geq 0$ by Proposition 4(iii), and $D_t \equiv Z_t^H - Z_t^D$. From the proof of Proposition 4, $D_t = (1 - \mu)D_{t-1} + e_{t-1}$

with $D_0 = 0$, so by forward substitution

$$D_t = \sum_{s=0}^{t-1} (1-\mu)^{t-1-s} e_s > 0 \qquad (t \geq 1). \tag{49}$$

Summability of e_t. Let Λ be a Lipschitz constant of f. Then $e_t \leq \varepsilon \Lambda \zeta \, |r_{t+1}^{old,H} - r_{t+1}^{old,D}|$, and by (18) each $|r_{t+1}^{old} - r^n| \leq C(\pi_t^{cp} + \pi_{t+1}^{cp})$ for a constant C depending on $(\phi_\pi, \psi, \phi_y, \kappa)$. Since $\pi_t^{cp} = \tilde{\pi}^{cp}(t/\hat{T}) e^{1-t/\hat{T}}$ satisfies $\sum_{t\geq 0} \pi_t^{cp} < \infty$ and $\pi_t^{cp} \to 0$, we obtain $e_t \to 0$ and $\sum_{t \geq 0} e_t < \infty$.

(a) $\mu \in (0,1)$. Fix $\eta > 0$ and choose S such that $e_s < \eta$ for all $s > S$. For $t > S + 1$,

$$D_t \leq (1-\mu)^{t-1-S} \sum_{s=0}^{S} e_s + \eta \sum_{j \geq 0} (1-\mu)^j = (1-\mu)^{t-1-S} \sum_{s=0}^{S} e_s + \frac{\eta}{\mu}, \tag{50}$$

so $\limsup_{t\to\infty} D_t \leq \eta/\mu$ for every $\eta > 0$, i.e. $D_t \to 0$; the convergence is geometric because e_t itself decays geometrically (up to a polynomial factor) and $(1-\mu)^t$ does. Yet $D_t > 0$ at every finite $t \geq 1$.

(b) $\mu = 0$. Then $D_t = \sum_{s=0}^{t-1} e_s$ is strictly increasing, bounded by $\sum_{s\geq 0} e_s < \infty$, and therefore converges to

$$\lim_{t\to\infty} D_t = \sum_{s=0}^{\infty} e_s = \varepsilon \sum_{s=0}^{\infty} \left[f(\rho_s^{g,H}) - f(\rho_s^{g,D}) \right] > 0, \tag{51}$$

which is strictly positive because every $e_s > 0$. No finite date after which the shock vanishes is used anywhere in the argument. □