CAPITAL GAINS TAXATION AND ASSET PRICE VOLATILITY*

PAU BELDA†

Abstract

Do capital gains tax cuts destabilize or stabilize asset prices? In an asset pricing model with heterogeneous agents and realization-based taxation, a tax cut has two opposing effects. It dampens volatility by reducing realization-based trading frictions, but also amplifies it by strengthening the pass-through from expectations to prices, fueling self-fulfilling fluctuations. Estimated on U.S. stock-market data, the model implies that the sequence of tax cuts since the 1970s triggered a net increase in volatility of about +35%, driven primarily by stronger belief-to-price pass-through. Policy experiments suggest a tax on unrealized gains robustly reduces volatility, whereas a financial transaction tax has mixed effects.

JEL codes: D83, D84, E44, G12, G14, H20, H31

*First version: July 2020. This version: February 2026. I am grateful to Albert Marcet for disciplining my thinking throughout the research process. I have also benefited from comments from Vladimir Asriyan, Jordi Caballé, Nicola Gennaioli, Dino Palazzo, Thomas Piketty, Ricardo Reis, Luis E. Rojas, Jaume Ventura, and audiences at numerous conferences and seminars. This version of the paper benefited from an extended research visit at CREi. I am grateful to my former colleagues in the Capital Markets Section of the Federal Reserve Board for stimulating conversations about the topic. The initial stages of the project were motivated by spirited discussions with José Luís Peña-Mir about the causes of the Great Depression. This research received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 788547, APMPAL-HET). Refine.ink was used to check the paper for consistency and clarity. The views expressed in this paper are those of the author and should not be interpreted as reflecting the views of the Bank of England.

†Research Economist at the Bank of England (Pau.Belda@bankofengland.co.uk).

1. Introduction

A longstanding question in public finance is whether a Capital Gains Tax (CGT) "*accentuates asset price fluctuations, promoting economic instability*" (Somers (1948)). The conventional view holds that realization-based taxes are destabilizing: by taxing gains upon sale, taxes may discourage investors from selling appreciated assets in booms and may accelerate the sale of depreciated assets in busts, amplifying price swings.[1] However, Haugen and Heins (1969) argue the opposite: by reducing after-tax expected returns, taxes can dampen the response of prices to shifts in expectations, thereby stabilizing markets. Which force dominates, and under what conditions, remains unresolved.[2]

Recent econometric evidence supports the stabilizing view. By exploiting variation in portfolios' tax exposure around the 1978 and 1997 U.S. tax reforms, Z. Dai, Shackelford, and Zhang (2013) show that cuts to the top long-term CGT increased stock return volatility. While this evidence identifies the net effect of taxes on volatility, distinguishing among competing mechanisms and quantifying their relative strength requires a structural model.[3]

However, incorporating realization-based capital gains taxation into equilibrium asset pricing models has proven challenging. Realization-based taxation introduces asymmetric trading frictions, path-dependence through cost bases and carryover losses, and the need for active trade among heterogeneous agents. These features have limited prior work. Analytical models that incorporate realization are either highly stylized or stop short of delivering explicit equilibrium pricing functions.[4] On the quantitative side, equilibrium models typically abstract from realization-based taxation, while models that include it focus on portfolio choice without equilibrium price determination.[5] As a result, we lack a dynamic equilibrium framework with realized-based taxation that can be used to quantify how capital gains taxation affects price volatility.

This paper addresses this gap by developing a general equilibrium asset pricing model with realization-based capital gains taxation that captures both channels. I first study an analytical version to characterize the theoretical mechanisms. Then, I estimate a quantitative version on U.S. stock market data to evaluate the effect of the tax cuts implemented in the US since the 1970s. The main finding is that tax cuts were destabilizing on net.

[1] This view dates to Somers (1948). Other early papers endorsing this view include Cloe (1952), Gemmill (1956), and Somers (1960).

[2] Haugen and Wichern (1973) argued that determining which effect dominates *"is a very difficult* [problem], *yet it must be considered important since the 'locked-in' issue has been used effectively as an argument against increasing the tax rate on capital gains.*"

[3] Z. Dai, Shackelford, and Zhang (2013) suggest another mechanism to rationalize the results: CGT as a risk-sharing device with the government. However, this hypothesis sits uneasily with the U.S. tax system, where tax liabilities depend on investors' realization decisions and loss offsets are severely constrained.

[4] Stiglitz (1983) analyzes a two-agent, two-state setup; Sahm (2008) considers a two-period environment; Klein (1999) and Viard (2000) derive Euler equations capturing lock-in but do not solve for equilibrium prices. Haugen and Wichern (1973) relies on simulations in a stylized setting with exogenous beliefs.

[5] See McGrattan and Prescott (2005), McGrattan (2023), Anagnostopoulos, Cárceles-Poveda, and Lin (2012), and Brun and González (2017) for the former; Ehling et al. (2010), M. Dai et al. (2015), and Ehling, Tompaidis, and C. Yang (2019) for the latter.

The analytical model is a microfounded version of Stiglitz (1983): two groups of investors with heterogeneous beliefs trade a Lucas tree subject to realization-based capital gains taxation. The economy alternates between boom and bust states, and agents agree to disagree about state persistence, generating trade.[6]

From an individual investor's perspective, capital gains taxation creates an inaction region. Selling triggers a tax on accrued gains, so an investor may value a position above its after-tax liquidation value and hold when she would otherwise sell. This no-trade zone expands with the tax rate—which is often referred to as a *lock-in effect*. In equilibrium, this lock-in tightens the conditions for trade: because sellers anticipate paying taxes upon realizing gains, marginal buyers in booms must be sufficiently optimistic to compensate sellers for the tax wedge. The tax thus acts as a selection device, making boom buyers more optimistic, which raises boom prices relative to bust prices. Moreover, as in Stiglitz (1983), for a given configuration of beliefs, the tax reduces the net payoff from a recovery, depressing bust prices relative to boom prices. The widening of the gap between boom and bust prices implies larger volatility.

Higher taxes, though, also reduce after-tax expected payoffs, attenuating the sensitivity of prices to belief revisions—which I refer to as a *pass-through effect*.[7] While Haugen and Heins (1969) anticipated this channel for exogenous beliefs, endogenizing them reveals how the pass-through effect depends critically on the nature of agents' uncertainty. In the two-state model, if agents know both the relevant stochastic processes and the equilibrium pricing function—the baseline Rational Expectations equilibrium in the paper—there is no effect. If agents instead receive noisy signals about fundamentals, the pass-through effect is static: taxes scale expected payoffs and hence scale price responses to belief changes. When agents must learn about the pricing function itself due to imperfect market knowledge—as in the Internal Rationality framework of Adam and Marcet (2011)—the effect is dynamically amplified: a higher tax dampens price responses to belief revisions, reduces forecast errors, induces smaller belief updates, and further stabilizes subsequent prices. A CGT thus weakens the price–belief feedback loop that generates excess volatility.

Overall, the net effect of capital gains taxation on price volatility is theoretically ambiguous; which effect dominates is an empirical question. To address it, I develop a quantitative version of the model that extends the analytical framework along several dimensions: a continuum of states, many investor types, a richer tax system featuring stochastic rates and carryover losses, and learning about both fundamentals and prices.

The model is computationally challenging, as it features heterogeneous agents, aggregate risk, subjective beliefs and learning, pronounced nonlinearities (including singularities), and several tax-related state variables. Building on Belda, Heineken, and Ifrim

[6] See Xiong (2013) for a review of the literature on heterogeneous beliefs.

[7] This mechanism is consistent with the micro-evidence in Giglio et al. (2021): equity shares are more aligned with investors' beliefs in tax-advantaged accounts, reflecting higher sensitivity of portfolios to beliefs when taxes are lower.

(2025), I tackle these challenges by exploiting a feature of the model's information setup. With Internal Rationality under imperfect market knowledge, individual rational agents take asset prices as a state variable, such that they do not need to track the entire cross-sectional distribution to solve their saving problems, circumventing what Moll (2025) labeled a "nonsensical" problem in heterogeneous-agent models.[8] The solution method decomposes naturally into two stages—solving the individual problem given idiosyncratic beliefs and computing market equilibrium—rendering the computational challenges tractable.

I estimate the full model by Simulated Method of Moments using U.S. stock-market data since 1975, a few years before the first statutory reduction in the capital gains tax rate in 1978. The model fits the targeted moments closely and also performs well on a set of validation moments excluded from the estimation. The moment set has two parts. The first targets standard objects in the excess volatility literature—means, volatilities, and comovements of prices, dividends, and valuation ratios, including a Campbell–Shiller decomposition of valuation dynamics and the response of valuations to price shocks.[9] The second targets capital gains realizations, in particular the tax elasticity of realizations that is central in the public-finance literature (e.g., Agersnap and Zidar (2021)). The model matches this elasticity at both medium- and long-run horizons, avoiding an implausibly large short-run trading response that would mechanically bias the estimated effects of taxation on volatility. Moreover, the model captures the correlation between price changes and trading volume well, as in Adam et al. (2015).

Leveraging the model's quantitative fit, I use counterfactual simulations to estimate the causal effect of the capital gains tax cuts implemented since 1975 on volatility and to provide an accounting of the mechanisms at work. The model implies that these tax cuts raised price–dividend volatility by about +35% relative to its observed level. The increase reflects two offsetting forces. Isolating the pass-through channel via a counterfactual without realization frictions, lower taxes raise volatility by approximately +50%. Realization-based trading frictions and their interactions with the pass-through partially offset this, accounting for the residual 15%. These findings are robust across tax-rate measures (effective and statutory) and across alternative informational assumptions about the tax process, including tax foresight and tax learning. The model-implied causal effect aligns with the external cross-sectional evidence in Z. Dai, Shackelford, and Zhang (2013), delivering an increase in stock return volatility of the same order of magnitude.[10]

[8] Examples of papers that solve quantitative heterogeneous models with nonlinearities and aggregate risk include Han, Y. Yang, et al. (2021), Kase, Melosi, and Rottner (2022), Fernández-Villaverde, Hurtado, and Nuno (2023) and Fernández-Villaverde (2025).

[9] Other papers using subjective expectations to explain excess volatility include Adam, Marcet, and Nicolini (2016), Adam, Marcet, and Beutel (2017), Barberis et al. (2018), and Bordalo et al. (2024).

[10] To quantitatively benchmark against Z. Dai, Shackelford, and Zhang (2013), who identify volatility effects around discrete reforms, I compute semi-elasticities within their event windows; the model delivers estimates of the same order of magnitude, providing a structural interpretation of their reduced-form findings.

Finally, I use the framework to evaluate alternative tax instruments. In the analytic model, an accrual-based capital gains tax retains the stabilizing pass-through effect while eliminating lock-in, since it does not distort trading decisions. By contrast, a Financial Transaction Tax (FTT) features both effects, leaving its net impact theoretically ambiguous—consistent with the mixed empirical findings in that literature.[11]

Quantitative experiments reinforce these conclusions. First, I implement a revenue-neutral regime change from realization-based to an accrual-based CGT system. Requiring an average rate of roughly 7%, the reform reduces valuation volatility by about 20%. Thus, a lower tax applied in a less distortionary manner appears to pay off in terms of stock market stability. Second, I introduce small supplementary taxes on either transactions or unrealized gains in the current realization-based regime, in line with existing policy proposals. A small FTT mainly lowers the persistence—and thus the long-run dispersion—of valuation ratios, while leaving short-run volatility largely unchanged. Return volatility responds weakly and non-monotonically, becoming destabilizing at higher rates. A small accrual-based tax on unrealized gains performs better: it delivers broad and monotonic stabilization by compressing short-run valuation fluctuations and attenuating shock propagation. Hence, from a financial stability perspective, supplementing the current system with a modest tax on unrealized gains appears to be the more reliable instrument.

The rest of the paper is organized as follows. Section 2 presents the analytic model and explores the mechanisms linking taxes and volatility. Section 3 sets up a quantitative version, presents an algorithm to solve and estimate it, and evaluates its performance. Section 4 computes the causal effect of tax cuts and presents a sensitivity analysis. Section 5 discusses a tax on unrealized gains and financial transactions. Section 6 concludes.

2. Analytic Model

This section develops the analytical framework. Section 2.1 lays out the economic environment: a two-state endowment economy with heterogeneous agents and realization-based capital gains taxation. Section 2.2 characterizes optimal portfolio choice and derives the no-trade region that gives rise to the lock-in effect. Section 2.3 introduces the equilibrium concepts used throughout the paper, ranging from Bayesian–Rational Expectations to Internal Rationality. Section 2.4 then studies how capital gains taxation affects price volatility across these environments, isolating the pass-through and lock-in channels and

[11] For instance, for Italy, Cappelletti, Guazzarotti, and Tommasino (2017) find that the FTT increased volatility, whereas Becchetti, Ferrari, and Trenta (2014) and gonzalez5062009revisiting report volatility declines for France and Spain, respectively; Eichfelder, Noack, and Noth (2022) finds only weak evidence of a persistent reduction in France. Using Chinese data, Deng, Liu, and Wei (2018) show that volatility falls primarily in markets with limited institutional participation. On the theory side, Buss and Dumas (2019) demonstrate that higher trading fees can raise return volatility, while Dávila (2023) emphasize that the asset-price impact is a priori indeterminate. Quantitative equilibrium analyses likewise suggest that FTTs need not reduce—and may even increase—volatility (e.g., Adam et al. (2015); Buss et al. (2016)).

characterizing conditions under which each dominates.

2.1. Model Structure

Consider a discrete-time economy where $t = 0, 1, 2, ...$ At each date, the economy can be in a high or low state, that is, $s_t \in \{H, L\}$. The sequence of states $\{s_t\}$ follows an exogenous time-homogeneous Markov chain governed by the transition probabilities

$$Pr(s_{t+1} = s' | s_t = s) = \pi_{s \to s'}$$

for a given s_0. The stationary distribution of the Markov chain is given by $Pr(s_t = H) = \mu$, $Pr(s_t = L) = 1 - \mu$.

There is a single risky asset in exogenous net supply normalized to 1 that pays $D_t = D(s_t)$ units of goods each period where

$$D(s_t) = \begin{cases} D_H & \text{if } s_t = H \\ D_L & \text{if } s_t = L \end{cases}$$

with $D_H > D_L$. At time zero, each investor owns S^i_{-1} shares, and a sufficiently large and non-state-contingent endowment W accrues in each period. Markets for the consumption good (the numeraire) and the stock are competitive.

Let $\tau \in [0, 1)$ be the capital gains tax rate, levied only on realized gains. The tax code reflects two main features of the real-world version: gains are only taxed when selling (realization-based), and the treatment of gains and losses is asymmetric (only net gains are taxed).

The budget constraint of any investor is given by

$$C^i_t + P_t S^i_t = W + (D_t + P_t) S^i_{t-1} - \tau \max\{0, S^i_{t-1} - S^i_t\} \max\{0, P_t - P^i_{bt}\} + T^i_t \qquad (1)$$

where P^i_{bt} is the basis price. The term $\tau \max\{0, S^i_{t-1} - S^i_t\} \max\{0, P_t - P^i_{bt}\}$ captures the fact that taxes are only paid when the asset is sold with positive gains. Agents are atomistic and take T^i_t as lump-sum. In equilibrium, taxes are transferred back to each individual according to her contribution to abstract from potential redistributive effects of τ which are not the focus of the paper. For the analytical version, I introduce two simplifications in the tax code. First, realized losses are not carried forward to offset future gains. Second, realized gains are measured as the one-period price change rather than against a historical basis, that is, $P^i_{bt} = P_{t-1}$. These assumptions streamline the exposition while preserving the mechanism of interest. Both features are reinstated in the quantitative analysis in Section 3.

Investor i faces the following savings-consumption problem

$$\max_{\{C_t^i, S_t^i \in \Gamma_t^i\}_{t=0}^{\infty}} \mathbb{E}_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t U(C_t^i) \tag{2}$$

where

$$\Gamma_t^i = \left\{ (C_t^i, S_t^i) : \text{budget constraint (1)}; \quad 0 \leq C_t^i; \quad 0 \leq S_t^i \leq \bar{S} \right\} \tag{3}$$

$\delta \in (0,1)$ is a discount factor. In this section, the utility function is assumed to be linear, $U(C) = C$. The feasibility set Γ includes the budget constraint (1) and limits on the choice variables. The lower bound on holdings represents a standard no-short selling constraint; the upper bound on the stock holdings helps to ensure the existence of a maximum, given risk neutrality.[12]

Agents optimize using a subjective probability measure $\mathcal{P}^i$, which sets the probability distributions about the variables beyond the investor's control, in line with the Internal Rationality framework of Adam and Marcet (2011). The underlying probability space is given by $(\Omega^i, \mathcal{B}^i, \mathcal{P}^i)$ where Ω is the sample space with $\omega = \{s_t, P_t, T_t^i\}_{t=0}^{\infty}$ as a typical element, $\mathcal{B}$ denotes the σ-algebra of Borel subsets of Ω and $\mathcal{P}$ is a subjective probability measure over $(\Omega, \mathcal{B})$. Note that prices and tax rebates are part of the sample space, as the tax rebate and the pricing functions are generally unknown.

Let the population have unit mass, split into two groups $\{A, B\}$ with measure $1/2$ each. Ex-ante they differ only in their subjective probability measure. In particular, they have different beliefs about the Markov transition probabilities. Investors in group A believe that the current state is more persistent: $\pi_{H\to H}^A > \pi_{H\to H}^B$ and $\pi_{L\to L}^A > \pi_{L\to L}^B$. Thus, type A investors are optimistic in good times but pessimistic in bad times.[13]

The subjective price probability distribution is discussed later; at this point, I only assume that it has bounded support and satisfies a Markov property. Investors take the tax rebate process $\{T_t^i\}$ as an exogenous Markov process independent of her trading decisions with bounded support.

2.2. Individual Optimal Demand

The investor's problem admits a recursive formulation. The state vector is

$$z_t^i = (s_t, P_t, P_{bt}^i, T_t^i, S_{t-1}^i), \tag{4}$$

[12] As shown by Stiglitz (1983) and Constantinides (1983), with frictionless shorting investors can hedge locked-in gains and replicate pre-tax allocations, rendering capital-gains taxation largely non-distortionary; prohibiting short sales makes the timing option non-tradeable, restores the lock-in effect, and lets the tax affect portfolios and prices.

[13] This belief heterogeneity can be microfounded through Bayesian learning with dogmatic priors. Investors observe state transitions to update beliefs about persistence, but with sufficiently strong priors, the weight on observed data vanishes and beliefs remain fixed at initial priors.

comprising the aggregate state s_t, the asset price P_t, the basis price P_{bt}^i, the transfer T_t^i, and the investor's lagged shareholdings S_{t-1}^i.[14] These states determine the resources of agent i at any point in time, being sufficient for the investor's decision problem. Note that z_t^i has a Markov structure under $\mathcal{P}^i$. The constraint correspondence Γ_t^i is non-empty, compact-valued, and continuous in z_t^i. Under the maintained assumptions of bounded support for prices, transfers, and endowments, the objective function is bounded.[15] The Bellman equation characterizing the investor's value function is

$$V^i(z_t^i) = \max_{(C_t^i, S_t^i) \in \Gamma_t^i} \left\{ C_t^i + \delta \, \mathbb{E}^{\mathcal{P}^i}[V^i(z_{t+1}^i) \mid z_t^i] \right\}. \tag{5}$$

Standard contraction mapping arguments establish the existence and uniqueness of V^i, along with an optimal policy correspondence characterized by the associated first-order conditions.

Lemma 1. *Define the investor's stock subjective marginal valuation as*

$$Q_t^i \equiv \delta \mathbb{E}^{\mathcal{P}^i} \left[\frac{\partial V^i(z_{t+1}^i)}{\partial S_t^i} \,\middle|\, z_t^i \right] \tag{6}$$

Then, the stock policy function is given by

$$S_t^i = \begin{cases} \bar{S} & \text{if } Q_t^i \geq P_t \\ 0 & \text{if } Q_t^i \leq P_t - \tau \max\{0, P_t - P_{bt}^i\} \\ S_{t-1}^i & \text{if } P_t - \tau \max\{0, P_t - P_{bt}^i\} \leq Q_t^i \leq P_t \end{cases} \tag{7}$$

Proof. Appendix D.1.

Figure 1 plots the resulting demand schedule under $P_t \geq P_{bt}^i$ and $Q_t^i > P_{bt}^i$, where P_{bt}^i is the basis price (in this section, it is fixed at P_{t-1}). Because the CGT applies only to realized gains, the investor's optimal position S_t^i is a three–region step function in the current price P_t. For prices

$$P_t \;\geq\; \frac{Q_t^i - \tau P_{bt}^i}{1 - \tau}$$

she finds the stock overvalued even in after–tax terms and optimally sells her entire position. For sufficiently low prices, $P_t \leq Q_t^i$, she views the stock as strictly undervalued and buys

[14] The treatment of prices as a state variable depends on investors' information. Under rational expectations with known pricing functions, P_t would be replaced by its determinants. When investors treat prices as exogenous, P_t enters directly as a state variable. I maintain the latter assumption for generality.

[15] Under $\mathcal{P}^i$, agents believe all the variables in the right-hand side of the budget constraint (1) have bounded support. Since $C_t \geq 0$ and $C_t \leq \bar{C}$ $\forall t$, where $\bar{C}$ is consumption in the most favorable combination of states, then $\mathbb{E}_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t C_t \leq \mathbb{E}_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t \bar{C} = \frac{\bar{C}}{1-\delta} < \infty$.

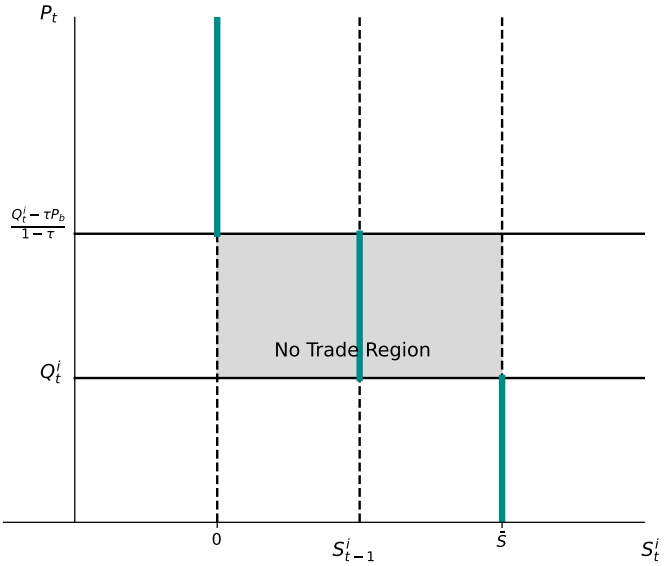


Figure 1: *Stock demand function. The graph plots the policy function in equation (7) for the case with $P_t \geq P_b$ and $Q_t^i > P_b$.*

up to the full position ($S_t^i = \bar{S}$). In between, the tax gives rise to an intermediate band,

$$Q_t^i \;\leq\; P_t \;\leq\; \frac{Q_t^i - \tau P_{bt}^i}{1-\tau}$$

in which the option value of deferring taxation on unrealized gains outweighs marginal incentives to trade, so she simply carries forward her previous holding ($S_t^i = S_{t-1}^i$). Thus, τ endogenously generates an inaction region between the buy and sell thresholds that prevents sales that would otherwise occur: this is the standard lock–in effect.

The width of the no-trade region is proportional to the valuation–basis wedge, $\frac{\tau}{1-\tau}\big(Q_t^i - P_{bt}^i\big)$. Holding the basis fixed, the inaction band widens when the investor's subjective valuation Q_t^i rises relative to P_{bt}^i. In equilibrium, valuations are typically higher in booms and the basis is relatively low compared with contemporaneous valuations, so the no-trade region tends to be larger precisely when embedded gains (measured at market prices) are large. Moreover, for a given $Q_t^i - P_{bt}^i$, the width is increasing in the tax rate. Hence a higher capital gains tax amplifies inaction exactly in states in which selling would otherwise help contain elevated prices, echoing classic lock-in arguments (e.g., Somers (1948)).

2.3. Equilibrium

The competitive equilibrium is standard: In each period, a price P_t and allocations $\{S_t^A, S_t^B\}$ must clear the market and satisfy the optimal policy (7) for each agent. As shown in the

proof of Lemma 1, the subjective marginal valuation of the stock is given by:

$$Q_t^i = \delta \mathbb{E}^{\mathcal{P}^i}\left[D_{t+1} + P_{t+1} - \tau \max\{0, P_{t+1} - P_t\}\mathbb{1}\{S_{t+1}^i < S_t^i\} \mid z_t^i\right] \tag{8}$$

which involves the expected selling probability. In general, this expectation is unknown, and numerical methods are needed. Instead, in this section, I construct an equilibrium that implies an analytically tractable form of this expected probability, which is illustrative of the role of CGT on price volatility.

To have a benchmark, let me briefly expose the equilibrium without taxes, which does not involve the expected probability of trading. In this case, the agent wants to hold everything if the stock is undervalued, hold nothing if the stock is overvalued, and hold anything if the stock is correctly priced, always with respect to her subjective valuation Q_t^i. With this demand schedule, market-clearing prices are such that the most optimistic agent is the one pricing and holding the asset. In what follows, I construct an equilibrium that preserves the trading strategies when taxes are introduced.

2.3.1. Bayesian Rational Expectations Equilibrium (B-REE)

In this section, I assume the subjective probability measures $\{\mathcal{P}^i\}$ are common knowledge. This knowledge about market participants and rationality (i.e., agents solve a well-posed optimization problem that implies that they value the asset using their Euler Equation) implies that agents can deduce the pricing function, that is, how beliefs and fundamentals map into prices. With this assumption, I construct an equilibrium in which ownership alternates with the state.[16]

Before stating the equilibrium, let me introduce some notation. Let $\mathbf{P}^* \equiv [\, P_H^* \quad P_L^* \,]'$ denote the vector of equilibrium state prices. Let

$$\mathbf{D}^e \equiv \begin{pmatrix} \pi_{H\to H}^A D_H + (1 - \pi_{H\to H}^A) D_L \\ \pi_{L\to L}^B D_L + (1 - \pi_{L\to L}^B) D_H \end{pmatrix} \equiv [\, D_H^e,\ D_L^e \,]'$$

be the vector of expected dividends in each state, evaluated under the beliefs of the marginal agent in that state. Let $S_t^j \equiv \int_0^{1/2} S_t^i \, di$ denote aggregate stock holdings of group $j \in \{A, B\}$.

To streamline the proposition statement below, define the upper bound on type B beliefs

$$\bar{\pi}_{H\to H}^B \equiv \frac{\left(1 - \pi_{L\to L}^B\right)\left(1 - \delta\, \pi_{H\to H}^A\right)}{1 - \delta\left(1 - \pi_{L\to L}^B\right)}$$

[16] I label the baseline as a Bayesian Rational Expectations Equilibrium (B-REE). Heterogeneous beliefs are rationalized as Bayesian updating from common public signals (prices) under dogmatic priors. Conditional on their posterior beliefs, agents understand the equilibrium pricing function and form model-consistent expectations of prices and dividends; in this sense the expectations are rational, even though beliefs differ across agents.

Next, define

$$M(\tau) \equiv (1-\delta\pi^A_{H\to H})\big(1-\delta[\pi^B_{L\to L}+\tau(1-\pi^B_{L\to L})]\big)-\delta^2(1-\pi^A_{H\to H})(1-\pi^B_{L\to L})(1-\tau) \quad (9)$$

and the matrix

$$K(\tau) \equiv \frac{\delta}{M(\tau)}\begin{pmatrix} 1-\delta[\pi^B_{L\to L}+\tau(1-\pi^B_{L\to L})] & \delta(1-\pi^A_{H\to H}) \\ \delta(1-\pi^B_{L\to L})(1-\tau) & 1-\delta\pi^A_{H\to H}\end{pmatrix}$$

Note that under the maintained parameter restrictions $\delta \in (0,1)$, $\pi^A_{H\to H}, \pi^B_{L\to L} \in [0,1]$, and $\tau \in [0,1]$, $M(\tau) > 0$ so that $K(\tau)$ is well-defined. With these objects in mind, consider the following proposition.

Proposition 1. *Assume that beliefs satisfy*

$$(i) \quad \pi^A_{H\to H} > 1-\pi^B_{L\to L} \quad (10)$$

$$(ii) \quad \pi^B_{H\to H} \leq \bar{\pi}^B_{H\to H} \quad (11)$$

Then, there exists a competitive equilibrium with:

$$(S^A_t, S^B_t) = \begin{cases}(1,0), & \text{if } s_t = H\\ (0,1), & \text{if } s_t = L\end{cases} \quad (12)$$

$$P_t = \begin{cases}P^*_H, & \text{if } s_t = H\\ P^*_L, & \text{if } s_t = L\end{cases} \quad (13)$$

where equilibrium state-prices are given by:

$$\mathbf{P}^* = K(\tau)\mathbf{D}^e \quad (14)$$

Proof. Appendix D.2.

Thus, there exists an equilibrium in which changes in the state of nature trigger trade that reallocates the asset to the group that is most optimistic in the current state. Absent taxes, this is the natural competitive allocation under heterogeneous beliefs, and any ordering $\pi^B_{H\to H} < \pi^A_{H\to H}$ suffices to sustain it. With realization-based capital gains taxation, however, selling triggers a tax liability that drives a wedge between the seller's reservation price and the market price, so market clearing requires strictly more disagreement than the baseline ordering.

Condition (i) ensures $P^*_H > P^*_L$, consistent with $D_H > D_L$: the marginal buyer in the boom must be sufficiently more optimistic than the marginal seller in the bust. Condition (ii) is the seller participation constraint: the belief gap $\pi^A_{H\to H} - \pi^B_{H\to H}$ must exceed the

wedge

$$\kappa(\tau) \;\equiv\; \frac{\tau\Big(\pi^{A}_{H\to H} - \big(1-\pi^{B}_{L\to L}\big)\Big)}{1-\delta\,\tau\big(1-\pi^{B}_{L\to L}\big)}$$

which is the minimum disagreement needed to induce type B holders to sell in the $L \to H$ transition despite incurring the realization tax. Because $\kappa(\tau)$ is strictly increasing in τ, higher taxes demand greater disagreement, tightening the admissible belief region. This requirement is implemented as an upper bound on $\pi^{B}_{H\to H}$: type B investors must be sufficiently pessimistic about the boom — pessimistic enough that the after-tax price still exceeds their valuation. Thus, lock-in operates here not by eliminating state-contingent trade, but by increasing the disagreement required to make trade privately optimal. If condition (ii) fails, the economy instead admits a locked-in no-trade equilibrium with degenerate allocations and price indeterminacy; Appendix B provides details.

More broadly, the belief restrictions reveal that realization-based taxation acts as a selection device on market participants. To sustain trade in the boom, sellers must hold sufficiently pessimistic beliefs (low $\pi^{B}_{H\to H}$) and buyers sufficiently optimistic ones (high $\pi^{A}_{H\to H}$), so the tax endogenously concentrates participation among agents with more polarized views.[17]

With this configuration of beliefs, equilibrium prices are pinned down by the Euler equations of the marginal traders, which imply a 2×2 linear system with solution (14). Prices are linear in expected dividends, with loading matrix $K(\tau)$ determined by belief heterogeneity and taxes. To lighten the notation, from now on I denote $\pi_A \equiv \pi^{A}_{H\to H}$ and $\pi_B \equiv \pi^{B}_{L\to L}$, unless otherwise stated.

2.3.2. B-REE with Learning about Dividends

This section relaxes the assumption of full information about the dividends process. In particular, agents only observe true dividends up to a noisy signal, $\mathcal{D}_t = D(s_t) + \varepsilon^{\mathcal{D}}_t$, with $\varepsilon^{\mathcal{D}}_t \sim \mathcal{N}(0, \sigma^2_{\varepsilon^{\mathcal{D}}})$. The noise might be due to measurement errors, imperfect information, or other frictions. To filter the noisy signal, agents use the following state-space model:

$$\mathcal{D}_t = H_t\,\mathbf{b}^{D}_t + \varepsilon^{D}_t \qquad \varepsilon^{D}_t \sim \mathcal{N}\big(0, \sigma^2_D\big) \tag{15}$$

$$\mathbf{b}^{D}_t = \mathbf{b}^{D}_{t-1} + \boldsymbol{\nu}^{D}_t \qquad \boldsymbol{\nu}_t \sim \mathcal{N}\big(\mathbf{0}, \sigma^2_{\nu^D}\mathbf{I}_2\big) \tag{16}$$

where $H_t \equiv \big(\mathbb{1}\{s_t = H\}, \mathbb{1}\{s_t = L\}\big)$ selects the realized Markov state $s_t \in \{H, L\}$, and $\mathbf{b}^{D}_t = (b^{D}_{H,t}, b^{D}_{L,t})'$ are the underlying dividend levels, which are thought to be very persistent. To estimate them, agents use a Kalman filter. Let $b^{D}_{st}|\mathcal{I}^{t-1} \sim \mathcal{N}(\beta^{D}_{st-1}, \sigma^2_{ds})$ be the prior belief based on information up to $t-1$ and $\sigma^2_{ds} = \frac{\sigma^2_{\nu^D} + \sqrt{\sigma^4_{\nu^D} + 4\pi_s\,\sigma^2_{\nu^D}\sigma^2_D}}{2\,\pi_s}$ is

[17] In a multi-agent setting, this logic implies that boom-state demand is endogenously concentrated among the most optimistic investors: the tax wedge raises the required valuation gap, pushing marginally optimistic buyers out of the market and leaving only the top tail of beliefs to absorb the asset.

the steady state forecast-error variance of the Kalman filter. Note it depends not only on the measurement precision $1/\sigma_D^2$ and on the volatility of the latent component $\sigma^2_{\nu^D}$, but also on the stationary probability with which state s is observed (e.g., $\pi_s = \mu$ for $s_t = H$). on measurement precision $1/\sigma_D^2$, process volatility $\sigma^2_{\nu^D}$, and the long-run visit probability π_s (e.g., $\pi_s = \mu$ for $s_t = H$), since more frequent observations leads to a smaller forecast-error variance. Optimal filtering implies that the posterior distribution when period t information is revealed is given by $b^D_{st}|\mathcal{I}^t \sim \mathcal{N}(\beta^D_{st}, \sigma^2_{ds})$. Then, steady-state optimal filtering implies

$$\beta^D_{st} = \beta^D_{st-1} + g_{ds} I_{st}\Big(\mathcal{D}_t - \beta^D_{st-1}\Big) \tag{17}$$

and

$$g_{ds} \equiv \frac{\sigma^2_{ds}}{\sigma^2_{ds} + \sigma^2_D} \tag{18}$$

where $I_{st} \equiv \mathbb{1}\{s_t = s\}$. Naturally, agents only update beliefs about the state s dividends when that state realizes and new dividends data is available. Otherwise, beliefs are not updated, that is, $\beta^D_{st} = \beta^D_{st-1}$. Altogether, the subjective expectations of state s dividends are given by $\mathbb{E}^{\mathcal{P}}_t\Big[D_{st+j}\Big] = \beta^D_{st}$.

The subjective price model nests Bayesian-Rational Expectations Equilibrium without Learning as a special case when agents assign probability 1 to $b^D_{s0} = D_s$ and believe $\sigma^2_{\nu^D} = 0$. Then, $\beta_{st} = D_s$ and prices are given by the B-REE prices in all periods.

In what follows, I consider small deviations from the B-REE, characterized by a small $\sigma^2_{\nu^D}$, such that the B-REE optimal trading strategies remain optimal. To focus on the effect of time-varying beliefs without adding an additional layer of heterogeneity, I assume that both agents hold the same dividend expectations. Under these assumptions, the system of Euler equations characterizing the B-REE becomes

$$P_{H,t} = \delta\Big[\pi_A\big(\beta^D_{H,t} + \mathbb{E}^{\mathcal{P}}_t P_{H,t+1}\big) + \big(1-\pi_A\big)\big(\beta^D_{L,t} + \mathbb{E}^{\mathcal{P}}_t P_{L,t+1}\big)\Big], \tag{19}$$

$$P_{L,t} = \delta\Big[\pi_B\big(\beta^D_{L,t} + \mathbb{E}^{\mathcal{P}}_t P_{L,t+1}\big) + \big(1-\pi_B\big)\Big(\beta^D_{H,t} + \mathbb{E}^{\mathcal{P}}_t P_{H,t+1} - \tau\big(\mathbb{E}^{\mathcal{P}}_t P_{H,t+1} - P_{L,t}\big)\Big)\Big]. \tag{20}$$

where the main difference with respect to the case without Learning is the time-varying nature of expected state-dividends and, as a result, the time-varying nature of prices. Proceeding by forward iteration and using a standard transversality condition, equilibrium prices become:

$$\boldsymbol{P}^*_t = K^D \boldsymbol{\beta}^D_t \tag{21}$$

where $\boldsymbol{P}^*_t = [P^*_{H,t}, P^*_{L,t}]$, $\boldsymbol{\beta}^D_t = [\beta^D_{H,t}, \beta^D_{L,t}]$, and

$$K^D \equiv \frac{\delta}{M^D}\,\mathcal{A}^D$$

with

$$M^D \equiv (1-\delta\pi_A)\left(1-\delta\tau(1-\pi_B)-\delta\pi_B\right)-\delta^2(1-\pi_A)(1-\pi_B)(1-\tau)$$

and

$$\mathcal{A}^D \equiv \begin{pmatrix} (1-\delta\tau(1-\pi_B))\pi_A+\delta\left(1-\pi_A-\pi_B\right) & \left(1-\delta\tau(1-\pi_B)\right)\left(1-\pi_A\right) \\ (1-\pi_B)\left(1-\delta\tau\,\pi_A\right) & \pi_B(1-\delta\pi_A)+\delta(1-\pi_B)(1-\tau)(1-\pi_A) \end{pmatrix}$$

Equilibrium prices (21) are well-defined provided $1-\delta\tau(1-\pi_B)>0$ and $M^D>0$, which always hold for the assumed bounds on the set of parameters. As in the case with full knowledge of the dividends process, equilibrium prices are a linear function of expected dividends. In this case, though, dividends beliefs are time-varying, which generates additional variability in prices.

2.3.3. Equilibrium with Learning about Prices

This section presents an alternative information setup that forces agents to learn from prices. This framework is interesting for two reasons: it captures the price extrapolation behavior frequently observed in markets, and it opens a new channel through which taxes affect volatility.

Agents have full knowledge of the dividend process but only imperfect knowledge of other agents' beliefs. This prevents them from computing the equilibrium pricing function—a calculation that would normally proceed by solving the system of Euler equations given knowledge of all agents' beliefs. Instead, agents must form probabilistic expectations about future prices directly, treating prices as uncertain even conditional on fundamentals. This departure from rational expectations equilibrium has significant implications for price dynamics and tax incidence.

This prevents them from computing the equilibrium pricing function—a calculation that would normally proceed by solving the system of Euler equations given knowledge of all agents' beliefs. Instead, agents must form probabilistic expectations about future prices directly, treating prices as uncertain even conditional on fundamentals.

In line with the noisy signal framework for dividends, let agents use the following stochastic model of prices:

$$P_t = H_t\,\mathbf{b}_t + \varepsilon_t^P \qquad \varepsilon_t^P \sim \mathcal{N}\big(0,\ \sigma_P^2\big) \tag{22}$$

$$\mathbf{b}_t = \mathbf{b}_{t-1} + \boldsymbol{\nu}_t \qquad \boldsymbol{\nu}_t \sim \mathcal{N}\big(\mathbf{0},\ \sigma_\nu^2\mathbf{I}_2\big) \tag{23}$$

where $H_t = \big(\mathbb{1}\{s_t=H\},\ \mathbb{1}\{s_t=L\}\big)$ is a 1×2 selector that picks out the realized state at date t and $\mathbf{b}_t = (b_{Ht}, b_{Lt})'$. Thus, agents think of prices as being made of a fundamental and very persistent component $\mathbf{b}$ plus noise. This model extends the standard linear model considered in the learning literature, and it is consistent with the extrapolative behavior

shown in surveys. The subjective price model nests Bayesian-Rational Expectations Equilibrium prices as a special case when agents assign probability 1 to $b_{s0} = P_s^*$ and believe $\sigma_\nu^2 = 0$. Then, $\beta_{st} = P_s^*$ in all periods, and prices are given by the B-REE prices in all periods.

Agents observe prices in state s only if that particular state is realized; otherwise, they do not update their beliefs. $\mathbf{b}$ is an unobserved component and ε_t^P and ν_t are innovations to the signal and the unobserved component, respectively. Since agents do not observe $\mathbf{b}$, they filter the signal using a Kalman filter, similar to the one used for learning about dividends. Optimal filtering implies that the posterior distribution when period t information is revealed is given by $b_{st}|\mathcal{I}^t \sim \mathcal{N}(\beta_{st}, \sigma_s^2)$.

As in Adam, Marcet, and Beutel (2017), I suppose the transitory price innovation decomposes into two independent pieces,

$$\varepsilon_t^P \;=\; \varepsilon_{t+1}^1 \;+\; \varepsilon_t^2 \tag{24}$$

where $\varepsilon_t^1 \sim i.i.d.\ \mathcal{N}(0, \sigma_b^2)$ and $\varepsilon_t^2 \sim i.i.d.\ \mathcal{N}(0, \sigma_p^2)$ are mutually independent, so that $\sigma_P^2 = \sigma_b^2 + \sigma_p^2$. Suppose agents observe ε_t^1 with a one-period lag, that is, at time t they receive the information shock $\varepsilon_t^b \;\equiv\; -\varepsilon_t^1$ with $\varepsilon_t^b \sim i.i.d.\ \mathcal{N}(0, \sigma_b^2)$, ε_t^b is measurable with respect to $\mathcal{I}^t$, and it is orthogonal to $\mathcal{I}^{t-1}$. Intuitively, ε_t^b captures the period-t news that helps agents disentangle the transitory component embedded in the previous price observation. Let $I_{st} \equiv \mathbb{1}\{s_t = s\}$. Then, under steady-state optimal filtering, posterior means update as

$$\beta_{st} \;=\; \beta_{s,t-1} \;+\; g_s\, I_{s,t-1}\Big(P_{t-1} - \beta_{s,t-1}\Big) \;+\; g_s\, I_{s,t-1}\varepsilon_t^b \tag{25}$$

where the steady-state Kalman gain is

$$g_s \;\equiv\; \frac{\sigma_s^2 + \sigma_\nu^2}{\sigma_s^2 + \sigma_\nu^2 + \sigma_P^2} \tag{26}$$

If the previous state differs from s, then $I_{s,t-1} = 0$ and beliefs are not updated, i.e. $\beta_{st} = \beta_{s,t-1}$. Altogether, subjective expectations satisfy $\mathbb{E}_t^{\mathcal{P}}[P_{s,t+1}] = \beta_{st}$.

In what follows, I consider small deviations from the B-REE with learning about dividends, characterized by a small σ_ν^2, such that the optimal B-REE trading strategies remain optimal. In this case, equilibrium prices satisfy

$$P_{H,t} = \delta\Big[\pi_A\big(D_H + \mathbb{E}_t^{\mathcal{P}} P_{H,t+1}\big) + \big(1-\pi_A\big)\big(D_L + \mathbb{E}_t^{\mathcal{P}} P_{L,t+1}\big)\Big] \tag{27}$$

$$P_{L,t} = \delta\Big[\pi_B\big(D_L + \mathbb{E}_t^{\mathcal{P}} P_{L,t+1}\big) + \big(1-\pi_B\big)\Big(D_H + \mathbb{E}_t^{\mathcal{P}} P_{H,t+1} - \tau\big(\mathbb{E}_t^{\mathcal{P}} P_{H,t+1} - P_{L,t}\big)\Big)\Big] \tag{28}$$

Using the subjective price beliefs and solving these equations, one obtains

$$\boldsymbol{P}_t^* = \delta\Big[\mathcal{A}\boldsymbol{D} + \mathcal{B}\boldsymbol{\beta}_t\Big] \tag{29}$$

where

$$\mathcal{A} = \begin{pmatrix} \pi_A & (1-\pi_A) \\ \dfrac{(1-\pi_B)}{1-\delta(1-\pi_B)\tau} & \dfrac{\pi_B}{1-\delta(1-\pi_B)\tau} \end{pmatrix}; \qquad \mathcal{B} = \begin{pmatrix} \pi_A & (1-\pi_A) \\ \dfrac{(1-\pi_B)(1-\tau)}{1-\delta(1-\pi_B)\tau} & \dfrac{\pi_B}{1-\delta(1-\pi_B)\tau} \end{pmatrix} \tag{30}$$

The pricing equation (29) and the updating equation (25) fully characterize the dynamics of the model when agents learn about prices. Instead of being equal to the present value of dividends, equilibrium prices depend on the expected one-period ahead dividends and expectations about one-period ahead prices. The previous models that assumed knowledge of the pricing function implied expected $t+1$ prices equal to expected future dividends, from period $t+2$ on. This formulation is then more general, allowing for price beliefs that deviates from the present value of dividends, in line with the empirical observations of excess volatility.

2.4. Taxes and Volatility

This section studies how capital gains taxation affects stock price volatility across the equilibrium environments introduced in Section 2.3. The analysis proceeds gradually to isolate the mechanisms at work. In the Bayesian–Rational Expectations Equilibrium, state prices are constant, so all price variation is between-state: volatility is entirely driven by the economy alternating between the high- and low-dividend states. In this benchmark, the capital gains tax amplifies volatility by widening the wedge between state-contingent prices. Introducing learning generates an additional source of within-state variation. Under the B-REE with dividend learning, time-varying dividend beliefs induce fluctuations in state prices even when the dividend state is fixed; in that case, the tax tends to reduce within-state volatility by dampening the pass-through from beliefs to prices. Finally, in the equilibrium with learning about prices, beliefs move endogenously with realized prices, strengthening within-state dynamics; capital gains taxes then reduce volatility through two margins, by lowering the beliefs-to-price pass-through and by attenuating belief fluctuations themselves. Thus, learning about prices strengthens the stabilizing role of capital gains taxes, but the same force is already present under rational expectations. Because taxation raises between-state dispersion but dampens within-state transmission, the net effect on overall volatility is theoretically ambiguous.

2.4.1. Bayesian-Rational Expectations

In the B-REE, prices fluctuate with fundamentals: as the economy alternates between recession and boom, prices move between P_L^* and P_H^*. The equilibrium pricing formula (14) shows that taxes are capitalized into prices, which may affect the price gap between states. The following proposition analyzes this (de)stabilizing role of the tax.

Proposition 2. *In the Bayesian-Rational Expectations Equilibrium, the variance of the stock price is increasing in τ:*

$$\frac{d}{d\tau} Var\Big(P_t^*\Big) > 0$$

Proof. Appendix D.3.

The proposition shows that, for the belief configuration that supports the B-REE, a realization-based capital gains tax widens the wedge between high-state and low-state prices. It is useful to separate two components of the total effect. The first holds beliefs fixed within the admissible region characterized in Proposition 1. The second considers how the tax shifts the boundary of admissible beliefs, and therefore changes equilibrium beliefs when the boundary binds.

Start with the fixed-belief component. Equilibrium prices solve the Euler system

$$P_H = \delta\big[\pi_A(D_H + P_H) + (1-\pi_A)(D_L + P_L)\big] \tag{31}$$

$$P_L = \delta\big[\pi_B(D_L + P_L) + (1-\pi_B)\big(D_H + P_H - \tau(P_H - P_L)\big)\big] \tag{32}$$

The tax enters in the low state because a holder in state L values a recovery partly through the associated capital gain $P_H - P_L$, of which a fraction τ must be paid in taxes upon realization. This reduces the perceived payoff from transitioning to H and therefore lowers P_L. The decline in P_L feeds back into the valuation in the high state through the first Euler equation, lowering P_H as well. The proof shows that P_L falls more than P_H: the former is hit directly, while the latter declines only indirectly through future expected prices. As a result, the price wedge widens and price volatility rises. This fixed-belief mechanism is a version of the result in Stiglitz (1983).

At the boundary, there is an additional effect. The tax rate also changes which belief configurations can sustain the equilibrium: a higher τ strengthens the lock-in force and makes it harder to induce realization, so sustaining trade requires more optimistic beliefs for agent A about boom persistence. This endogenous belief adjustment further increases P_H more than P_L, widening the price wedge. Therefore, when the boundary binds, the same tax change widens the wedge both directly through realization taxation and indirectly through belief selection, amplifying the increase in volatility relative to the fixed-belief benchmark.[18]

[18] The boundary adjustment $\pi_A'(\tau) > 0$ is derived via the implicit function theorem applied to the seller

2.4.2. B-REE with Learning about Dividends

The learning process through which agents cope with their imperfect knowledge of the dividend process generates variation in state prices. Whereas previously agents knew with certainty the dividend level associated with each state, under noisy information they must learn about the true underlying level. For instance, observing a higher-than-expected dividend makes them more optimistic about the underlying level, while a lower-than-expected dividend prompts a downward revision in beliefs. These fluctuations in beliefs create additional volatility that was absent under full information. The new source of volatility opens a new channel for taxes to affect volatility.

Remember $\underline{\pi}_A$ and $\bar{\pi}_A$, the lower and upper bounds on π_A stated in Proposition 1. Define the auxiliary upper bound

$$\bar{\pi}_A^{\text{aux}}(\pi_B) \;\equiv\; \begin{cases} \dfrac{\delta(1-\pi_B)}{\delta(2-\pi_B)-1} & \text{if } \; \delta(2-\pi_B) > 1 \\ 1 & \text{if } \; \delta(2-\pi_B) \leq 1 \end{cases}$$

and let

$$\hat{\pi}_A \;\equiv\; \min\{\bar{\pi}_A,\; \bar{\pi}_A^{\text{aux}}(\pi_B)\}$$

Consider the set

$$\mathcal{I}^D(\pi_B) \;\equiv\; \Big\{\pi_A \in [\underline{\pi}_A, \hat{\pi}_A] :\; F_H(\pi_A;\pi_B) > 0 \text{ and } F_L(\pi_A;\pi_B) > 0\Big\}$$

where, for $s \in \{H, L\}$,

$$F_s(\pi_A;\pi_B) \;\equiv\; (1-\delta\pi_A)\, q(\pi_A)\, \Sigma_\beta^D\, a_s(1)' \;-\; a_s(0)\, \Sigma_\beta^D\, a_s(0)'$$

Here $\Sigma_\beta^D \equiv \text{Var}(\boldsymbol{\beta}_t^D)$, $q(\pi_A) \equiv (\pi_A, 1-\pi_A)$, and $a_s(\tau)$ denotes the s-th row of the tax-dependent dividend-learning loading matrix $\mathcal{A}^D(\tau)$. In particular, $a_s(0)$ and $a_s(1)$ are the same loadings evaluated at $\tau = 0$ and $\tau = 1$, respectively. Given these thresholds, the following lemma formalizes the effect of taxes on the volatility of each state-price.

Lemma 2. *Consider the Bayesian–Rational Expectations Equilibrium with dividend learning under the mean–coefficient approximation with $\lambda_s \equiv \mathbb{E}[I_{s,t}]$ and symmetric Kalman gains $g_H = g_L \equiv g > 0$.*[19] *Suppose that $\pi_A \in \mathcal{I}^D(\pi_B)$. Then, for all $\tau \in (0,1)$ and for each $s \in \{H, L\}$,*

$$\frac{d}{d\tau}\, Var(P_{s,t}^*) \;<\; 0$$

participation constraint, which requires $\tau(1+\delta(1-\pi_B)) < 1$ for the boundary to be locally well-defined. This condition is satisfied for all empirically relevant tax rates (e.g., at $\delta = 0.95$ and $\pi_B = 0.8$ it permits τ up to 0.84, well above any observed capital gains tax rate). For this reason, I do not elevate it to a formal assumption in Proposition 2.

[19] Note that in the two-state case $\lambda_H = \mu$ and $\lambda_L = 1-\mu$.

Proof. Appendix D.4.

In the dividend-learning equilibrium, a higher capital gains tax rate dampens price volatility because it reduces the sensitivity of equilibrium valuations to the dividend-belief vector through the pricing loadings. In our specification, belief dynamics are pinned down by dividend realizations (and hence are unaffected by τ); taxes operate solely through the equilibrium mapping from beliefs to prices.

At the same time, taxes affect the common scale factor in the pricing kernel, $\delta/M^D(\tau)$, which mechanically rescales state prices. Holding fixed the normalized dispersion of the state-specific loadings, an increase in $\delta/M^D(\tau)$ can therefore raise $\mathrm{Var}(P^*_{s,t})$. The net effect of taxes on volatility reflects the interaction between (i) a variance-reducing pass-through channel operating through the loadings and (ii) an offsetting variance-increasing rescaling through $\delta/M^D(\tau)$.

The belief restriction in Lemma 2 guarantees that the pass-through attenuation dominates the countervailing level rescaling uniformly over $\tau \in (0,1)$ in both states. Formally, $F_s(\pi_A;\pi_B) \geq 0$ is a variance-dominance condition under the weighting matrix Σ^D_β: it ensures that the contraction in the state-s pricing loadings induced by taxation is strong enough to offset the mechanical variance increase arising from the common scale factor $\delta/M^D(\tau)$.

Economically, the restriction rules out belief configurations in which prices effectively load on only one component of the dividend-belief state—either because $q(\pi_A)$ is close to a corner or because the relevant loading vector $a_s(\cdot)$ places negligible weight on one belief component. In such near-degenerate cases, taxes compress risk exposure primarily along a single dimension, and this compression can be too weak (relative to the change in $\delta/M^D(\tau)$) to deliver a monotone decline in price volatility. By restricting attention to $\mathcal{I}^D(\pi_B)$, the lemma ensures that both belief components enter state prices with nontrivial weight, so the attenuation in pass-through is sufficiently strong to dominate the level-rescaling effect. The condition is therefore best interpreted as a transparent sufficient condition for monotone volatility effects, rather than as a necessary requirement.

2.4.3. Learning about prices

Consider the equilibrium with learning about prices. The key insight is that the pass-through effect compounds dynamically: a lower pass-through makes prices less reactive to shocks, which implies less variation in the forecast errors and more stable beliefs. This section formalizes this point.

A key difference with respect to Learning about dividends is that belief dynamics

become self-fulfilling. In particular, they are characterized by the following VAR(1):[20]

$$\boldsymbol{\beta}_t = \boldsymbol{d}_t + \boldsymbol{U}_t \boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t \tag{33}$$

with

$$\begin{aligned} \boldsymbol{d}_t &= \delta\, G\, \mathrm{diag}(\mathbf{I}_{t-1})\, \mathcal{A}\, \boldsymbol{D} \\ \mathbf{U}_t &= I_2 - G\, \mathrm{diag}(\mathbf{I}_{t-1}) + \delta\, G\, \mathrm{diag}(\mathbf{I}_{t-1})\, \mathcal{B} \\ \boldsymbol{\eta}_t &= G \mathbf{I}_{t-1} \varepsilon_t^b \end{aligned}$$

and $G = \mathrm{diag}(g_H, g_L)$, $\mathbf{I}_{t-1} = (I_{H,t-1}, I_{L,t-1})'$. This self-referential process generates an additional channel for taxes. To see that, take the variance on both sides of the pricing function to get this decomposition:

$$\mathrm{Var}(P_{st}) = \delta^2 \mathcal{B}_s \Sigma_\beta \mathcal{B}_s' \tag{34}$$

where $\Sigma_\beta \equiv \mathrm{Var}(\boldsymbol{\beta}_t)$ and $\mathcal{B}_s$ refer to the corresponding row of matrix $\mathcal{B}$. Its total derivative with respect to τ reads as

$$\frac{d}{d\tau}\mathrm{Var}(P_{st}) = 2\delta^2 \left(\frac{d\mathcal{B}_s}{d\tau}\right) \Sigma_\beta\, \mathcal{B}_s' + \delta^2\, \mathcal{B}_s\, \frac{d\Sigma_\beta}{d\tau}\, \mathcal{B}_s' \tag{35}$$

The first term captures the loadings effect (how τ changes $\mathcal{B}_s$), similar to the case with Learning about dividends; the second term captures how τ changes the beliefs' covariance matrix Σ_β. While with Learning about dividends, $\frac{d\Sigma_{\beta D}}{d\tau} = 0$ (entrywise), that is no longer true when agents learn about prices. The following lemma derives the conditions under which both terms in equation (35) are negative.

Before stating the result, introduce compact notation for the parameter restrictions used in the proof. Define the variance ratio

$$\rho \equiv \frac{\mathrm{Var}(\beta_{Lt})}{\mathrm{Var}(\beta_{Ht})} \geq 0$$

For given (δ, τ) and a candidate equilibrium variance ratio ρ, define

$$\underline{\pi}_B \equiv 1 - \frac{1}{\delta(2-\tau)}$$

[20] The VAR results from plugging the equilibrium pricing function (29) into the updating equation (25) for each state, using that in equilibrium agents observe the market-clearing price $P_{t-1} = P^*_{t-1}$. While agents may entertain a subjective signal extraction problem in which prices contain transitory components, the law of motion in (33) describes equilibrium belief dynamics under realized prices; hence ε_t^b should be interpreted as a belief (or information) shock.

and

$$\bar{\pi}_B(\rho) \;\equiv\; \frac{(1-\tau)(2\delta-1)+\sqrt{(1-\tau)^2(2\delta-1)^2+4\delta\big((1-\tau)+\rho\big)(1-\tau)(1-\delta)}}{2\delta\big((1-\tau)+\rho\big)}$$

Define the threshold

$$\Theta \;\equiv\; \max\big\{\Theta_1,\;\Theta_2\big\}$$

where Θ_1 and Θ_2 are given by

$$\Theta_1 \;=\; \frac{1-\delta\big(\pi_B+\tau(1-\pi_B)\big)}{\Xi}$$

$$\begin{aligned}\Xi = 1-\delta\pi_A+\delta^2(\pi_A-\pi_B)(\pi_A+\pi_B-1)\\ +\tau\,\delta(\pi_B-1)\Big[1+\delta(\pi_B-1)+\delta^2(\pi_A-1)(\pi_A+\pi_B-1)\Big]\end{aligned}$$

and

$$\Theta_2 \;=\; \frac{\delta\,\pi_B\Big[\pi_A-(1-\tau)(1-\pi_B)\Big]+\Big(1-2\pi_B\Big)-\tau(1-\pi_B)}{(1-\delta\pi_A)(\pi_A+\pi_B-1)\big(1-\delta(1-\pi_B)\tau\big)}$$

Given this notation, consider the following lemma.

Lemma 3. *Consider the equilibrium with learning about prices and a mean–coefficient approximation of equation (33) with $\mathbb{E}(I_{st})=\lambda_s$. Assume symmetric constant gains $g_H = g_L \equiv g > 0$. Suppose that*

$$\begin{aligned}&(i)\quad \pi_B \in \big(\underline{\pi}_B, \bar{\pi}_B(\rho)\big) \qquad \text{and} \qquad \bar{\pi}_B(\rho) \;>\; \underline{\pi}_B\\ &(ii)\quad \frac{\lambda_H}{\lambda_L} \;\geq\; \Theta\end{aligned}$$

Then there exists $\bar{g} > 0$ such that, for all $g \in (0, \bar{g})$ and each $s \in \{H, L\}$:

1. *The tax weakly reduces the pricing loading component,*

$$\Big(\frac{\partial \mathcal{B}_s}{\partial \tau}\Big)\Sigma_\beta\,\mathcal{B}_s' \;\leq\; 0$$

2. *The tax strictly reduces the belief dispersion component,*

$$\mathcal{B}_s\,\frac{d\Sigma_\beta}{d\tau}\,\mathcal{B}_s' \;<\; 0$$

Hence, the tax strictly reduces the volatility of equilibrium prices in state s,

$$\frac{d}{d\tau}\,Var\big(P_{st}^*\big) \;<\; 0$$

Proof. Appendix D.5.

Let me elaborate on the intuition behind the lemma. In the low state, the tax affects the price in two different ways, as apparent by inspecting the loadings matrix (30). A tax drags down the after-tax gain in the event of a recovery from L to H, so the low-state price becomes less sensitive to recovery beliefs. This is the stabilizing bit: shocks to expectations about future high prices translate less into today's low price, forecast errors shrink, belief updating becomes more cautious, and the whole price–belief feedback loop is weakened. However, the tax also introduces a basis-price channel: whenever realized gains are positive, the after-tax payoff loads on the basis with coefficient proportional to τ, so basis fluctuations become more price-relevant as τ rises. In the one-period-basis version used here, this channel is particularly transparent because the basis at $t+1$ equals P_{Lt}, which makes the wedge appear in the denominator and creates a direct feedback from the low-state price into the future tax term. This effect can be destabilizing by amplifying the impact of basis movements, even as the numerator channel tends to dampen within-state fluctuations. If fluctuations in beliefs are more relevant than fluctuations in basis prices for understanding price volatility, the stabilizing effect would prevail.

The restrictions on beliefs ensure that the stabilizing effect dominates. Assumption (i) is mild in the applications of interest because the variance ratio ρ is typically close to, or below, one, so the admissible interval $(\underline{\pi}_B, \bar{\pi}_B(\rho))$ is non-empty; the restriction simply rules out parameterizations in which low-state belief variation is implausibly large relative to the high state. When π_B is too small, the basis price dominates due to the nonlinear effect; when π_B is too large, the low state is so persistent that recoveries are almost never expected, so the stabilizing effect essentially disappears. The tax reduces volatility by weakening the sensitivity of the low-state price to the belief factor β_H. The restriction $\lambda_H/\lambda_L \geq \Theta$ is a sufficient dominance condition that ensures that the stationary covariance of beliefs is not overwhelmingly driven by low-state episodes in a way that could offset the tax-induced compression of the low-state loading on β_H. Under this condition, the reduction in the low-state exposure to β_H translates into an unambiguous decline in unconditional price volatility as τ rises. The small-gain assumption plays a complementary role: it linearizes the learning dynamics so that beliefs follow a time-invariant VAR(1) with a well-defined stationary covariance, and price fluctuations remain in a narrow band around their state-dependent means. A small g is consistent with the small deviation from the B-REE we are considering and with standard estimates in the literature (see Section 3).

The restrictions on beliefs are exactly what ensure that the stabilizing effect dominates. When π_B is too small, the basis price dominates due to the nonlinear effect; when π_B is very large, the low state is so persistent that recoveries are almost never expected, so the stabilizing effect essentially disappears. The condition on λ_H/λ_L ensures that a large share of stationary belief dispersion comes from fluctuations in β_H rather than from rare

low-state episodes. Since τ reduces volatility by compressing the low-state price's loading on β_H, this restriction makes that channel first-order for overall volatility. The small gain assumption plays a complementary role: it linearizes the learning dynamics so that beliefs follow a time-invariant VAR(1) with a well-defined stationary covariance, and price fluctuations remain in a narrow band around their state-dependent means. A small g is consistent with the small deviation from the B-REE we are considering and with standard estimates in the literature (see Section 3).

2.4.4. A more general proposition

This section integrates the results from the previous ones. In Proposition 2, state prices were constant and taxes influenced the gap between boom and bust prices; on the contrary, Lemmas 2 and 3 focused on price variation within each state. The following proposition integrates both rationales.

Proposition 3. *In the learning equilibrium linearized around the Bayesian–Rational Expectations Equilibrium, and under the conditions of Proposition 2 and Lemmas 2 and 3, the total effect of a capital gains tax on price volatility is ambiguous. In particular, to first order around the B-REE,*

$$\frac{d\ Var(P_t)}{d\tau} \approx \mu(1-\mu)\underbrace{\frac{d}{d\tau}\left[P_H^* - P_L^*\right]^2}_{>0} + \underbrace{\frac{d}{d\tau}\Big[\mu\ Var(P_{Ht}) + (1-\mu)\ Var(P_{Lt})\Big]}_{<0} \tag{36}$$

where P_s^* denotes the B-REE state price and $Var(P_{st})$ is computed under the stationary distribution of the linearized learning dynamics within state $s \in \{H, L\}$.

Proof. *The proposition follows from the law of total variance*

$$\mathrm{Var}(P_t) = \mathrm{Var}\big(\mathbb{E}[P_t \mid s_t]\big) + \mathbb{E}\big[\mathrm{Var}(P_t \mid s_t)\big]$$

along with the combination of Proposition 2 and Lemmas 2 and 3. Proposition 2 carries over to the learning economy because, under the mean-coefficient first-order approximation around the B-REE, the unconditional mean of state prices coincides with the B-REE price up, that is $\mathbb{E}(P_{st}) \approx P_s^*$ (under the assumptions of both Lemmas). ■

The capital gains tax has an ambiguous effect on price volatility in models with information frictions, contrasting with the unambiguous destabilizing effect under the Bayesian-Rational Expectations equilibrium. This ambiguity arises from the decomposition of total variance into two distinct components with opposing responses to taxation. The first component—the variance of conditional expectations across states—behaves like the B-REE case (up to a linear approximation) and then the Proposition 2 goes through. The second

component—the expected conditional variance—captures within-state price fluctuations generated by changes in expected state-payoffs and becomes stronger when agents extrapolate from prices.[21]

2.4.5. A remark

A natural concern is that the variance results above reflect a mechanical level effect: if taxes lower mean prices, a decline in $\text{Var}(P_t)$ could simply mirror that decline rather than a genuine reduction in price dispersion. The economically relevant object is $\text{Var}(\ln P_t)$, which nets out mean-valuation shifts.

The analytical results extend naturally to log-variance in some cases. For Proposition 2 the extension is immediate; for Lemma 2 it requires an additional condition on D_H/D_L that is straightforward to verify.[22] For Lemma 3, however, the extension does not follow from the same approach. In Lemma 2, the belief mean $\mathbb{E}[\beta_t^D] = D$ τ-independent, so the log transformation introduces a mean-level effect that can be controlled by a condition on dividends. In Lemma 3, the self-referential structure of learning about prices makes $\mathbb{E}[\beta_t]$ track the equilibrium price level $P_s^*(\tau)$, which depends on τ even under symmetric dividends. This introduces a mean-level effect of the same order as the variance compression that cannot be signed analytically. For this reason, I state the theoretical results in terms of $\text{Var}(P_t)$ and use the quantitative model to verify that the identified mechanisms are strong enough to reduce log-variance as well.

Moreover, the analytical section works with prices while the quantitative model targets the price-dividend ratio. This difference is purely one of normalization. In the analytical model, dividends are stationary, so prices inherit stationarity and are the natural object of analysis. In the quantitative model, dividends grow, so prices are nonstationary and the price-dividend ratio becomes the appropriate stationary variable. The theoretical results of Section 2 apply to whichever object is stationary in each setting: prices in the simple economy, the log price-dividend ratio in the richer one.

3. Quantitative Model

3.1 Structure

The backbone of the quantitative model is the model presented in Section 2. In this section, I extend it in key dimensions to make it closer to the real-world tax system and more suitable for quantitative analysis. In what follows, I describe all the extensions.

[21] In the model, the two variance components map neatly into fundamental and non-fundamental forces. However, all that is needed for the pass-through effect to work is variation in the expected state-payoffs. Setups with information frictions are a natural way of generating such variation, but the rationale is more general.

[22] These derivations are available upon request.

Stochastic capital taxes. The quantitative model incorporates a richer tax environment along several dimensions. A proportional tax on dividends τ_t^D is introduced, completing the characterization of personal capital income taxation. Moreover, I depart from the assumption of constant tax rates by allowing both dividend and capital gains taxes to evolve stochastically over time. Tax rates follow unit root processes:

$$\tau_t^i = \tau_{t-1}^i + \varepsilon_t^{\tau^i} \quad \text{for } i \in \{D, CG\} \tag{37}$$

where the innovations are jointly normal and potentially correlated:

$$\begin{pmatrix} \varepsilon_t^{\tau^D} \\ \varepsilon_t^{\tau^{CG}} \end{pmatrix} \sim \text{i.i.d.} \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\tau^D}^2 & \sigma_{D,CG} \\ \sigma_{D,CG} & \sigma_{\tau^{CG}}^2 \end{pmatrix} \right) \tag{38}$$

The unit root specification captures the high persistence of tax policy while allowing for permanent shifts such as major tax reforms. The correlation structure $\sigma_{D,CG}$ reflects the possibility that dividend and capital gains tax rates move together, for instance following joint tax reforms (e.g., the Economic Growth and Tax Relief Reconciliation Act of 2001).

Stochastic taxes introduce an additional source of volatility: changes in taxes will directly affect expected payoffs, demands and prices. Thus, while the theory in Section was about tax levels, here tax changes are a direct source of volatility themselves. When evaluating tax cuts, I will disentangle the effect of pure tax volatility from the effect of the tax level on volatility.

Carryover losses. In the previous sections, realized losses did not originate either a rebate when realized or a right to offset future gains. Thus, these "lost" losses are the polar opposite of full tax rebates which are often assumed in the literature. In this section, I relax this assumption, making it closer to the existing tax system: losses can be used to offset future gains, thus reducing future tax liabilities. This use of losses modifies the investor's budget constraint as follows:

$$\begin{aligned} C_t^i + P_t S_t^i =& W_t + ((1-\tau_t^D)D_t + P_t)S_{t-1}^i \\ & - \tau_t^{CG} \max\left\{ \max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i) - L_t^i S_{t-1}^i, 0 \right\} + T_t^i \end{aligned} \tag{39}$$

L_t^i are losses per share such that total carryover losses $L_t^i S_{t-1}^i$ are subtracted from gross realized gains, $\max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i)$, where P_{bt}^i is the basis price. Thus, the presence of carryover losses generates a second kink in the budget constraint at $\max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i) = L_t^i S_{t-1}^i$. Hence, the agent can sell $S_{t-1}^i - S_t^i = \min\left\{ S_{t-1}^i, \frac{L_t^i S_{t-1}^i}{(P_t - P_{bt}^i)} \right\}$ shares with embedded gains (i.e. $P_t > P_{bt}^i$), but does not pay taxes, as carryover losses exactly offset the gains at that point.

The stock of unrealized losses per share becomes a state variable. Its law of motion is

$$L^i_{t+1} = \begin{cases} \dfrac{\max\left\{L^i_t S^i_{t-1} - (P_t - P^i_{bt})\max\{0, S^i_{t-1} - S^i_t\},\, 0\right\}}{S^i_t} & \text{if } S^i_t > 0, \\ 0 & \text{if } S^i_t = 0 \end{cases} \tag{40}$$

Realized gains $(P_t - P^i_{bt}) > 0$ draw down losses up to the zero lower bound, while realized losses $(P_t - P^i_{bt}) < 0$ augment it. If no shares are sold, losses simply carry forward unchanged, i.e. $L^i_{t+1} S^i_t = L^i_t S^i_{t-1}$ for $S^i_t > 0$. At full liquidation ($S^i_t = 0$), the per-share loss state is set to zero and thus does not affect decisions off the market.[23]

Basis price. In this section, I allow the basis price to differ from the previous period price. Potentially, investors' portfolio contains shares acquired at different prices across multiple periods. To account for this heterogeneity, I track the average price basis using a quantity-weighted recursion similar to the literature (e.g., Ehling, Tompaidis, and C. Yang (2019)). When the agent increases her position from S^i_{t-1} to S^i_t, the new average basis reflects both the existing shares carried at basis P^i_{bt} and the newly acquired shares purchased at the current price P_t:

$$P^i_{bt+1} = \begin{cases} \frac{P^i_{bt} S^i_{t-1} + P_t (S^i_t - S^i_{t-1})}{S^i_t} & \text{if } S^i_t > S^i_{t-1} \\ P^i_{bt} & \text{if } S^i_t \leq S^i_{t-1} \end{cases} \tag{41}$$

This recursion preserves the total cost basis of the portfolio: the numerator represents the sum of the historical cost of existing shares and the cost of new purchases, while the denominator normalizes by the new total position. When the agent reduces her position or holds it constant, the average basis remains unchanged.

Risk aversion. The quantitative model includes a CRRA utility function with parameter γ. In this case, diminishing marginal utility ensures that the first-order conditions pin down interior allocations rather than the bang-bang solutions obtained under linearity.

Discount factor and supply shocks. The quantitative model incorporates stochastic variation in both the discount factor and aggregate stock supply. Specifically,

$$\tilde{\delta}_t = \delta + \varepsilon^\delta_t \tag{42}$$

and

$$\tilde{S}_t = 1 + \varepsilon^s_t \tag{43}$$

where innovations are i.i.d. with zero mean. Discount factor shocks capture dynamics in

[23] An alternative is to take the total unrealized losses $\mathcal{L}^i_t \equiv L^i_t S^i_{t-1}$ as the state variable and allow for losses survival even when equity positions have been liquidated using this law of motion $\mathcal{L}^i_{t+1} = \max\left\{\mathcal{L}^i_t - (P_t - P^i_{bt})\max(0, S^i_{t-1} - S^i_t),\, 0\right\}$.

the risk-free rate despite the absence of an explicit bond market in the model. Supply shocks account for time variation in the effective stock supply driven by share repurchases.

A continuum of states of nature. The baseline two-state economy is replaced by a richer environment with a continuum of states. While the analytical model restricted the economy to binary states $s_t \in \{H, L\}$ with corresponding discrete dividend and wage realizations, I now specify fundamental processes that follow continuous distributions (following Adam, Marcet, and Beutel (2017)):

$$\begin{aligned} \ln\frac{D_t}{D_{t-1}} &= \ln\beta^D + \ln\varepsilon_t^D \\ \ln\Big(1+\frac{W_t}{D_t}\Big) &= (1-p)\ln(1+w) + p\ln\Big(1+\frac{W_{t-1}}{D_{t-1}}\Big) + \ln\varepsilon_t^W \end{aligned} \tag{44}$$

Dividend growth follows a log-normal process with mean growth rate β^D and innovation ε_t^D, capturing both growth and fluctuations in stocks' payoffs. The wage-dividend ratio follows a persistent AR(1) process in logs with persistence parameter p. The intercept $\ln(1+w)$ pins down the deterministic steady state of $\ln(1+W_t/D_t)$. Wage and dividends innovations are jointly Normal and potentially correlated (capturing, for instance, aggregate income shocks):

$$\begin{pmatrix} \ln\varepsilon_t^D \\ \ln\varepsilon_t^W \end{pmatrix} \sim ii\mathcal{N}\left(-\frac{1}{2}\begin{pmatrix} \sigma_D^2 \\ \sigma_W^2 \end{pmatrix}, \begin{pmatrix} \sigma_D^2 & \sigma_{DW} \\ \sigma_{DW} & \sigma_W^2 \end{pmatrix} \right) \tag{45}$$

This extension enables quantitative analysis by matching empirically observed moments of income dynamics (e.g., dividends growth volatility). Even more importantly, the extension introduces dividend growth, which induces corresponding price appreciation over time. Growth overcomes a critical limitation of the two-state model, that only generates capital gains in the transition from the low to the high state, ignoring the possibility of accumulating gains over time as prices grow (e.g., prices in low states today might be much higher than in high states 30 years ago). This mechanical constraint severely limited the model's ability to match empirical patterns of capital gains realization.

A large population of heterogeneous groups with imperfect information. Instead of 2 types of agents, I consider a large finite number of I types with mass $1/I$ each, which allows for richer trading dynamics. As in the analytic model, agents differ in their subjective probability measure $\mathcal{P}^i$. Agents are endowed with perfect knowledge of the stochastic processes for wages and taxes, but they have imperfect knowledge about the stochastic process for dividends and the equilibrium pricing function (as in Section 2.3.3.). Investors deal with this friction by learning about both dividends and price growth.

In particular, agents do not know the mean dividend growth β^D. Instead, they think that dividends growth is the sum of two unobserved components, one persistent and the

other transitory:

$$\begin{aligned} \ln \frac{D_t}{D_{t-1}} &= \ln d_t^i + \ln \varepsilon_t^{d,i} \\ \ln d_t^i &= \ln d_{t-1}^i + \ln \nu_t^{d,i} \end{aligned} \tag{46}$$

where $\ln \varepsilon_t^{d,i} \sim ii\mathcal{N}(-\frac{(\sigma_d^2)^i}{2}, (\sigma_d^2)^i)$ and $\ln \nu_t^{d,i} \sim ii\mathcal{N}(-\frac{(\sigma_\nu^2)^i}{2}, (\sigma_\nu^2)^i)$. Signal extraction using the steady state Kalman filter implies that the posterior for the unobserved state is given by $\ln d_t | \mathcal{I}_t \sim \mathcal{N}(\ln(\beta_t^d)^i, (\sigma_{\beta_D}^2)^i)$. The posterior variance is given by $(\sigma_{\beta^d}^2)^i = (\sigma_d^2)^i \frac{-\lambda^i + \sqrt{(\lambda^i)^2 + 4\lambda^i}}{2}$ where $\lambda^i \equiv \frac{(\sigma_\nu^2)^i}{(\sigma_d^2)^i}$ is the signal-to-noise ratio. In turn, the posterior mean evolves as

$$\ln(\beta_t^d)^i = \ln(\beta_{t-1}^d)^i + g^i \left(\ln \frac{D_t}{D_{t-1}} - \ln(\beta_{t-1}^d)^i \right) + \kappa_d^i \tag{47}$$

where $\kappa_d^i \equiv g^i \frac{(\sigma_d^2)^i}{2} - (1 - g^i) \frac{(\sigma_\nu^2)^i}{2}$, and the steady-state Kalman gain is given by $g^i \equiv \frac{(\sigma_{\beta_D}^2)^i + (\sigma_\nu^2)^i}{(\sigma_{\beta_D}^2)^i + (\sigma_\nu^2)^i + (\sigma_d^2)^i}$. Altogether, expected dividend growth for agent i is given by

$$\mathbb{E}_t^{\mathcal{P}^i} \left[\frac{D_{t+1}}{D_t} \right] = \kappa^i (\beta_t^d)^i \tag{48}$$

where $\kappa^i \equiv \exp\{(\sigma_{\beta_d}^2)^i / 2\}$ is a variance correction term. This process implies that agents with a higher Kalman gain, who are more confident about the information content of dividend signals, are more optimistic when good news arrives and more pessimistic when bad news arrives, in line with the switching optimism of Section 2. Thus, agents agree-to-disagree, based on the different views they hold about the information content of dividend data.

Furthermore, the assumption of homogeneous subjective price beliefs is relaxed. That introduces another layer of disagreement and a trading motive based on disagreement about non-fundamentals. The subjective model for prices is analogous to the dividend-belief model in the sense that beliefs are updated by a forecast-error correction (Kalman-gain) rule. The timing, however, differs because the information set for prices follows Section 2.3.3: agents observe the market-clearing price P_t when they trade at t, but they only learn part of the transitory innovation embedded in the previous price observation with a one-period lag. Consequently, the signal that is informative about price growth and available at time t is lagged price growth, $\ln(P_{t-1}/P_{t-2})$ (together with the period-t information shock that helps disentangle the transitory component of P_{t-1} as in Section 2.3.3).

Formally, β_t^i is predetermined at the start of period t (measurable with respect to the information set $\mathcal{I}_t^i$ available before the period-t trading decision), and it governs agent i's conditional forecast of next period price growth. The updating equation is

$$\ln \beta_t^i = \ln \beta_{t-1}^i + g^i \left(\ln \frac{P_{t-1}}{P_{t-2}} - \ln \beta_{t-1}^i \right) + \kappa_p^i \tag{49}$$

with $\kappa_p^i \equiv g^i \frac{(\sigma_p^2)^i}{2} - (1 - g^i) \frac{(\sigma_{\nu p}^2)^i}{2}$.[24] Hence,

$$\mathbb{E}_t^{\mathcal{P}^i} \left[\frac{P_{t+1}}{P_t} \right] = \tilde{\kappa}^i \beta_t^i \tag{50}$$

where $\tilde{\kappa}^i \equiv \exp\{(\sigma_\beta^2)^i/2\}$, and β_t^i are the mean posterior beliefs about price growth.[25] Thus, the learning rules for dividends and prices share the same adaptive-learning structure, while the use of lagged price growth reflects the assumed lag in observing part of the transitory price shock.

3.2. Equilibrium computation

In this section, I characterize equilibrium in the extended model. I begin by restating the recursive problem of an individual investor. The Bellman equation of agent's i solves

$$V^i(z_t^i) = \max_{C_t^i, S_t^i} \left\{ \frac{(C_t^i)^{1-\gamma}}{1-\gamma} + \delta \, \mathbb{E}^{\mathcal{P}^i}[V^i(z_{t+1}^i) \mid z_t^i] \right\} \tag{51}$$

subject to the budget constraint (39) and stock and consumption limits.

In the extended economy, the vector of states for individual i becomes

$$z_t^i = \{W_t, D_t, \tau_t^{CG}, \tau_t^D, T_t^i, S_{t-1}^i, P_{bt}^i, L_t^i, P_t, (\beta_t^D)^i, \beta_t^i, \tilde{\delta}_t, \tilde{S}_t\}$$

which includes wages, dividends, taxes and transfers, past holdings, the basis price, carry-over losses per share, the stock price, beliefs on dividends and price growth, the discount factor, and aggregate supply.[26] Relative to the analytical benchmark, the state space is larger because it must accommodate the richer tax system, belief heterogeneity, and aggregate shocks. I use the following equilibrium concept.

Definition 1 (Recursive Competitive Equilibrium under Internal Rationality)**.** *Given the exogenous processes, agents' subjective probability measures* $\{\mathcal{P}^i\}_{i=1}^I$*, and initial individual states, a recursive competitive equilibrium under internal rationality is a collection of:*

[24] As in Adam and Merkel (2019), in the quantitative version, I set the information about the lagged transitory component equal to zero; otherwise, the updating rule would include an additional term $g^i \ln \varepsilon_t^b$, as in the analytic specification. Consistent with the timing conventions in Section 2.3.3, the price signal enters with a one-period lag through $\ln(P_{t-1}/P_{t-2})$, so β_t^i is predetermined at the start of period t.

[25] Adam, Marcet, and Nicolini (2016) show that, in closely related stock-pricing learning environments, discrepancies between the perceived and the actual price law of motion can be empirically hard to detect in aggregate data, so the learning rule need not imply large or easily detectable systematic forecast errors.

[26] Transfers are taken as given by individual households (small-agent assumption). In the quantitative model, I use either no transfers (baseline) or a rebate rule such that T_t^i is predetermined at the start of period t observed before the trading decision; see the sensitivity analysis is Section 4.2.

1. *Individual policy functions for consumption and stock holdings,*

$$\begin{pmatrix} C_t^i \\ S_t^i \end{pmatrix} = \boldsymbol{h}^i\big(z_t^i\big) \tag{52}$$

2. *A pricing function,*

$$P_t^* = p\big(\Gamma_t, W_t, D_t, \tau_t^{CG}, \tau_t^D, \tilde{\delta}_t, \tilde{S}_t\big) \tag{53}$$

where Γ_t is a cumulative distribution function over predetermined individual states,

$$\Gamma_t(s, b, l, \beta^D, \beta, \mathcal{T}) \equiv \sum_{i=1}^{I} \mu^i \mathbf{1}\Big\{S_{t-1}^i \leq s,\ P_{bt}^i \leq b,\ L_t^i \leq l,\ (\beta_t^D)^i \leq \beta^D,\ \beta_t^i \leq \beta,\ T_t^i \leq \mathcal{T}\Big\} \tag{54}$$

Such that:

(a) *Household optimality. For each agent i, the policy $\boldsymbol{h}^i$ solves the Bellman problem (51), subject to the period budget constraint (39), the asset-holding limits, and the subjective law of motion for individual states and beliefs implied by $\mathcal{P}^i$.*

(b) *Market clearing.*[27] *For all t,*

$$\sum_{i=1}^{I} \mu^i h_s^i(z_t^i) = \tilde{S}_t \tag{55}$$

$$\sum_{i=1}^{I} \mu^i C_t^i = D_t + W_t \tag{56}$$

Note that in conventional models with Rational Expectations, households forecast future prices by forecasting the aggregate state. Formally, one specifies a transition law for the cross-sectional distribution, $\Gamma_{t+1} = H(\Gamma_t, \cdot)$, and a pricing function $p(\Gamma_t, \cdot)$, so that agents form expectations about $p(\Gamma_{t+1}, \cdot)$ using H. Under internal rationality, an equilibrium transition law for Γ_t still exists because individual states evolve endogenously and markets clear each period. The key difference is informational: households do not know the mapping $p(\Gamma, \cdot)$. They therefore treat the current price as an individual state variable and forecast future prices directly using their subjective price model.

This equilibrium concept has a direct computational implication. Under each agent's subjective probability measure $\mathcal{P}^i$, the price process is an exogenous component of the perceived individual state, so the agent's dynamic program can be solved conditional on simulated price paths drawn from $\mathcal{P}^i$. Given the resulting policy rules, the equilibrium price

[27] Condition 56 assumes that taxes are rebate to investors. If one assumes taxes are thrown away, the condition becomes $\sum_{i=1}^{I} \mu^i C_t^i = D_t(1 - \tau_t^D) + W_t - \tau_t^{CG} \sum_i \mu^i \max\Big\{\max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i) - L_t^i S_{t-1}^i, 0\Big\}$.

in each period is obtained by aggregating individual demands and imposing stock-market clearing.[28] Operationally, computation proceeds in two steps: (i) solve the individual problems conditional on the perceived price process to obtain policy functions; (ii) use these policies to compute the market-clearing price and thereby recover the implied equilibrium pricing function $p(\Gamma, \cdot)$.[29]

Step 1: Recursive solution of the agent's consumption-savings problem.

For each agent, a recursive solution to her optimization problem boils down to a pair of unknown time-invariant policy functions for goods and stock demands $(C_t^i)^* = h_c^i(z_t^i)$, $(S_t^i)^* = h_s^i(z_t^i)$. Key challenges in approximating the policy functions include the presence of a highly non-linear subjective conditional expectation in the optimality conditions, the piecewise optimality conditions with a no-trade region, and the potentially binding stock holding limits. In what follows, I describe the strategy to tackle these challenges.

Denote $\psi(z_t^i) \equiv \frac{1}{P_t}\delta \mathbb{E}^{\mathcal{P}^i}\left[\frac{\partial V^i(z_{t+1}^i)}{\partial S_t^i} \,\middle|\, z_t^i\right] = \frac{Q_t^i}{P_t}$. Thus, $\psi(z_t^i)$ summarizes investor i's discounted subjective marginal value of carrying one additional share into $t+1$ relative to the current price. Under realization-based capital gains taxation with average-cost basis and loss carryovers, an additional share at t affects next period's direct tax liabilities and also the next-period tax states $(L_{t+1}^i, P_{b,t+1}^i)$. Appendix E derives the particular form of $\psi(z_t^i)$.

Under risk neutrality, the relevant comparison is simply between Q_t^i and the market price P_t (penciling in the corresponding tax payments). Under risk aversion, instead, the continuation value is weighted by marginal utility, so the optimality conditions can hold at an interior portfolio choice and consumption enters through $u'(C_t^i) = (C_t^i)^{-\gamma}$. In this case, the KKT conditions read as:

[28] I do not provide a general uniqueness result for the market-clearing price. In the simulations reported below, the numerical procedure did not encounter multiplicity.

[29] This two-step organization does not remove the endogenous interaction between beliefs and prices. At the beginning of each period t, individual states z_t^i are predetermined, including beliefs $(\beta_t^D)^i$ and β_t^i. Given a candidate price P, households' optimal policies $h_s^i(\cdot)$ imply demands $S_t^i(P)$. The equilibrium price P_t is pinned down by stock-market clearing, $\sum_i \mu_i S_t^i(P_t) = \widetilde{S}_t$. This defines the equilibrium pricing function $P_t^* = p(\Gamma_t, \cdot)$ as the market-clearing mapping from the current cross-sectional distribution and aggregate states into prices. Along the simulated equilibrium path, the realized market-clearing price P_t^* enters next-period individual states and belief updating, so beliefs affect prices through demand, and prices affect beliefs through updating.

$$
\begin{aligned}
C_t^i \geq \psi_t^{-1/\gamma} \quad &\text{and} \quad S_t^i = \bar{S} \\
C_t^i = \psi_t^{-1/\gamma} \quad &\text{and} \quad S_t^i \in [S_{t-1}^i, \bar{S}] \cup [\hat{S}_t^i, S_{t-1}^i] \\
\psi_t^{-1/\gamma} \geq C_t^i \geq \left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \quad &\text{and} \quad S_t^i = \hat{S}_t^i \ \text{ if } \ P_t > P_{bt}^i \\
C_t^i = \left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \quad &\text{and} \quad S_t^i \in [0, \hat{S}_t^i] \\
\left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \geq C_t^i \quad &\text{and} \quad S_t^i = 0
\end{aligned}
\tag{57}
$$

where $\psi_t \equiv \psi(z_t^i)$ and $\kappa_t^i \equiv 1 - \tau_t^{CG} \max\left\{\left(1 - \frac{P_{bt}^i}{P_t}\right), 0\right\}$, which captures the wedge between the market price and the after-tax liquidation value when unrealized gains are positive (and satisfies $\kappa_t^i = 1$ when $P_t \leq P_{bt}^i$). $\hat{S}_t^i$ denote the loss-exhaustion kink for the period-t choice S_t^i, i.e., the smallest post-trade position at t such that carryover losses exactly offset realized gains at t. Equivalently, when $P_t > P_{bt}^i$,

$$
\hat{S}_t^i = S_{t-1}^i \max\left\{0,\ 1 - \frac{L_t^i}{P_t - P_{bt}^i}\right\}
$$

and $\hat{S}_t^i = 0$ otherwise. Selling is tax-free for $S_t^i \in [\hat{S}_t^i, S_{t-1}^i]$ and taxable for $S_t^i < \hat{S}_t^i$. In the continuation-value term, the analogous next-period threshold $\hat{S}_{t+1}^i$ enters through the event $\{S_{t+1}^i < \hat{S}_{t+1}^i\}$. See Appendix E for details.

Conditions (57) characterize the investor's optimal stock demand by comparing the continuation value ψ_t^i to the marginal cost of adjusting the position. When the investor increases her holdings ($S_t^i > S_{t-1}^i$), the marginal cost is the purchase price P_t, and the interior FOC takes the usual form $(C_t^i)^{-\gamma} = \psi_t^i$. When the investor reduces her holdings, sales are tax-free as long as carryover losses offset realized gains, so the relevant marginal proceeds are P_t for $S_t^i \in [\hat{S}_t^i, S_{t-1}^i]$, and the same interior condition applies. Selling beyond the loss-exhaustion threshold ($S_t^i < \hat{S}_t^i$) triggers capital-gains taxation, so marginal proceeds fall to $\kappa_t^i P_t$ and the selling FOC becomes $(C_t^i)^{-\gamma}\kappa_t^i = \psi_t^i$. Because $P_t > \kappa_t^i P_t$ whenever $\kappa_t^i < 1$, the kink at $S_t^i = \hat{S}_t^i$ generates an inaction condition: if ψ_t^i lies between the two marginal-utility-scaled prices, the KKT inequalities hold and the optimum is $S_t^i = \hat{S}_t^i$ (and reduces to $S_t^i = S_{t-1}^i$ when $\hat{S}_t^i = S_{t-1}^i$). Finally, complementary slackness with the bounds $S_t^i \in [0, \bar{S}]$ delivers corner solutions, while interior choices satisfy the equalities above.

There are two main challenges to solve this problem: the unknown conditional expectation and the inequality constraints. The conditional expectation is parametrically approximated following the idea of the Parameterized Expectations Algorithm (Marcet (1988)); inequality constraints are dealt in a case-by-base basis building on the algorithm in Marcet and Singleton (1999). Appendix E contains the details.

Step 2: Aggregation and market-clearing prices

Once the first stage has been completed, equilibrium prices are determined by the stock market clearing condition:

$$\sum_i \mu^i S(z_t^i) = \tilde{S}_t \tag{58}$$

where $\mu^i = 1/I$ is the population mass of each group of investor. All it takes is to solve this nonlinear equation for P_t.

Solving the stock market clearing condition (58) for prices, one obtains that the equilibrium pricing function depends on the cross-sectional distribution and the macroeconomic variables:

$$P_t^* = p(\Gamma_t, W_t, D_t, \tau_t^{CG}, \tau_t^D) \tag{59}$$

where Γ_t is the current cumulative distribution over stock holdings, basis price, beliefs, and transfers:

$$\Gamma_t \equiv \Gamma(s, b, \beta^D, \beta, \mathcal{T}) = \sum_{i=1}^{I} \mathbf{1}\Big\{S_{t-1}^i \leq s, P_{bt}^i \leq b, (\beta_t^D)^i \leq \beta^D, \beta_t^i \leq \beta, T_t^i \leq \mathcal{T}\Big\} \tag{60}$$

Thus, with imperfect knowledge about the pricing function, Γ_t is needed to compute aggregate endogenous variables like prices, but not for solving the investors' problems.

3.3. Model parameterization

This section estimates the model using U.S. stock market data. The moments include standard unconditional statistics from the excess-volatility literature as well as conditional moments related to capital gains realization behavior. Because, despite the remarkable simplification from Internal Rationality, solving the model remains computationally intensive, I follow Adam, Marcet, and Beutel (2017) and estimate only a subset of parameters via SMM, calibrating the rest to direct data counterparts or to standard values in the literature. I first describe the data moments that discipline the estimation, then turn to the calibration choices and the SMM procedure.

3.3.1. Data moments

The moments fall into two groups: standard statistics capturing excess volatility in stock prices, and measures of how investors' realization behavior responds to capital gains taxes. The data sample is from 1975 to 2022, at quarterly frequency where available and annual otherwise. The asset pricing data focuses on the S&P500.

Facts about the US stock market. The statistics in Table 1 shows the classic "excess volatility" facts. The volatility of log returns is large, $\sigma(\ln R_t) = 15.00$, while dividend

Table 1: Stock market statistics. 1975-2022. *This table reports a list of statistics using the data sources described in Appendix A. The variables are in real terms. Growth rates and returns are annualized and multiplied by 100. Correlations are multiplied by 100. Newey-West standard errors are in parentheses, except for the IRFs that are obtained via bootstrapping with 1000 repetitions.* $\bar{d}_t = \sum_{j=1}^{\infty} \hat{\rho}^{j-1} \Delta ln D_{t+j}$ *and* $\bar{r}_t = \sum_{j=1}^{\infty} \hat{\rho}^{j-1} \ln R_{t+j}$ *come from a standard Campbell-Shiller decomposition.*

	Excess Volatility	Statistic	St. Errors
Variance of the PD ratio	$\mathbb{V}(\ln PD_t)$	19.12	3.57
Covariance PD - future dividends	$\mathbb{C}(\ln PD_t, \bar{d}_t)$	4.60	0.84
Covariance PD - future returns	$\mathbb{C}(\ln PD_t, \bar{r}_t)$	-14.52	2.78
Std. Deviation of log returns	$\sigma(\ln R_t)$	15.00	1.37
Std. Deviation of log dividends growth	$\sigma(\ln D_t/D_{t-1})$	3.40	0.50
Serial correlation PD ratio	$\text{corr}(PD_t, PD_{t-1})$	98.12	4.35
	Valuations, returns and capital gains		
Mean Price-Dividend ratio	$\mathbb{E}(PD_t)$	44.74	3.74
Mean stock returns	$\mathbb{E}(R_t)$	8.97	2.33
Mean dividend growth	$\mathbb{E}(D_t/D_{t-1} - 1)$	2.50	1.07
Mean stock price growth	$\mathbb{E}(P_t/P_{t-1} - 1)$	6.13	2.29
	Shock propagation		
Valuations response	$\mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 1, I_t) - \mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 0, I_t)$	3.21	0.38
	Trading Volume		
Comovement trading and prices	$\text{corr}(TV_t, \|P_t/P_{t-1} - 1\|)$	32.30	12.66

growth is comparatively smooth, $\sigma(\Delta \ln D_t) = 3.40$. Valuation ratios move a great deal: the dispersion of $\ln PD_t$ is 19.12 while its mean is just 3.80 in ln terms. The Campbell–Shiller decomposition reveals that most movements in valuations are accounted for by future returns rather than future cash flows. While this is broadly in line with classic results, the cash-flow share in this sample is 25%—notably higher than in Cochrane (2009)—and its sign is positive. In terms of levels, average returns are dominated by capital gains rather than dividend growth (6.13% versus 2.50%), underscoring that dividends capture only a relatively small portion of total payoffs during this period.

The table also reports a complementary measure of valuation volatility based on impulse responses: the one-year-ahead response of the log price–dividend ratio to a unit price-growth shock is 3.21, indicating a strong propagation of short-run shocks into valuations. Finally, following Adam et al. (2015), I include the correlation between detrended trading volume and absolute price changes.[30] While trading volume appears to be uncorrelated with valuations when trends are removed, periods of large price changes are significantly associated with high trading volume.

Capital Gains Tax Reforms. Since the 1970s, there have been significant changes in the statutory rates. For reference, Appendix C offers a brief summary. However, statutory rates might, in practice, be very different from effective marginal rates, which are the

[30]The table removes a linear trend, but results are robust to other detrending methods.

ones relevant for making decisions. Thus, the literature on capital taxes uses the so-called effective average marginal rates, that is, a value-weighted mean of the marginal tax rates of investors adjusting for the features of the tax code (as maximum and minimum taxes, partial inclusion of social security or phaseouts of the standard deduction) (see, for instance, McGrattan and Prescott (2005)). The NBER-TaxSim program computes these rates.

There is an important feature the NBER-TaxSim rates do not consider: capital gains accruing to pension funds, individual retirement accounts, and nonprofit organizations which are subject to a tax-deferred scheme.[31] Following the literature (e.g., McGrattan and Prescott (2003), Sialm (2009)), I also adjust the rates for the non-taxable share of equity income. I computed it following Rosenthal and Austin (2016). This adjustment is relevant due to the important regulatory changes involving pensions savings vehicles took place along the period (see McGrattan and Prescott (2005) for details). The result of these regulatory reforms was an enormous switch in corporate equities ownership, moving from households (taxable units) to pensions funds and IRA accounts (non-taxable units). See Appendix F for details in the computation. Figure 2 plots the top statutory rate and the effective average marginal tax rate. As observed, the effective rate is similar to scaling down the top statutory rate.

The tax elasticity of realization. A key debate in the capital gains taxation literature is the tax elasticity of realization, i.e., how much agents change realization behavior in response to a tax change.

In a recent paper, Agersnap and Zidar (2021) use a panel design that exploits cross-state policy variation and fixed effects while controlling for pre- and post-tax dynamics to isolate the effects of a particular tax change.[32] Instead, aiming at applying a comparable design to the simulated aggregate series generated by the model, I estimate the following

[31] The tax-deferred scheme implies that these funds are only taxed upon withdrawal. In this case, they face an ordinary income tax. As a result, they do not affect stocks equilibrium valuations. To see that, take a Traditional IRA, for instance. Contributions to it are deductible (under some conditions), which means that the investor is paying lower taxes (\$1 invested in the equity fund costs $\$(1-\tau)$). So, she saves taxes today with her contributions. She gets some income once she withdraw her money, which ends being $\$(1-\tau)$ per each dollar withdrawn. Hence, these future taxes are already offset by today taxes and have nothing to do with asset pricing. In fact, the tax deduction can imply an income tax cut since when the investor is working she usually is in a higher income bracket than when she gets retired, but this is another topic.

[32] Agersnap and Zidar (2021) implement a dynamic event-study with (i) a long-window regressor, $\Delta_{11} \log n_{s,t}$, and (ii) a stack of $\pm$3-year binned leads/lags of $\Delta_3 \log n_{s,t+r}$ to net out anticipatory and post effects, leveraging cross-state tax variation and rich fixed effects. Their medium-run elasticity is formed as a difference of coefficients, $\beta_\ell - \tilde{\beta}_{-3}$, from two complementary projections. However, in a single national series, those overlapping bins are nearly collinear with the long-window change, leaving little residual variation and producing unstable (often very large) coefficients. Equation (61) therefore uses a parsimonious long-difference design without the bin stack to target the same horizon directly and avoid near-collinearity in aggregate data.

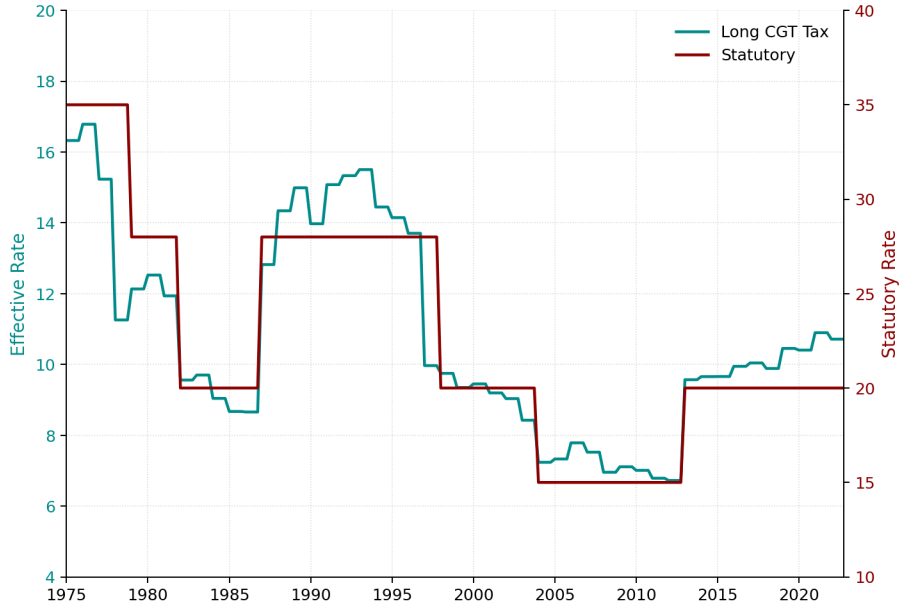


Figure 2: Statutory vs. Effective Long Capital Gains Tax Rate.

long-difference regression:

$$\Delta_n \log G_{t+h} = \alpha_\tau + \beta_\tau \, \Delta_n \log(1 - \tau_t^{CG}) \; + \; X_t \Lambda \; + \; u_t, \tag{61}$$

where G_t are aggregate realized gains, $\Delta_n x_t \equiv x_t - x_{t-n}$, and X_t is a vector of controls (all in log differences).

This specification relies on medium-run changes in aggregate time series. Under the assumption that over n years no omitted aggregate factor is both strongly trending and tightly correlated with $\Delta_n \log(1-\tau_t^{CG})$ conditional on X_t, β_τ in (61) is a consistent estimate of the medium-run net-of-tax elasticity of realizations. It can be converted to a tax-rate elasticity noting that

$$\frac{d \log(1-\tau_t)}{d \log \tau_t} = \frac{d \log(1-\tau_t)}{d\tau_t} \cdot \tau_t = -\frac{\tau_t}{1-\tau_t},$$

so the tax-rate elasticity evaluated at a benchmark τ^* is

$$\varepsilon_{\text{tax}}(\tau^*) = \beta_\tau \Big(- \tfrac{\tau^*}{1-\tau^*} \Big). \tag{62}$$

I include two groups of controls to purge economy-wide forces that comove with realizations but are not the capital-gains tax itself. First, I include a small set of minimal controls: the PD ratio, which absorbs valuation cycles that comove with realizations, and the net-of-dividend tax rate, to net out concurrent reforms that alter the relative price of payout modes and hence the propensity to realize gains. The second group includes

standard macro-financial variables: real GDP growth and the output gap, inflation, the federal funds rate, the three-month T-bill rate, and the ten-year Treasury yield. These controls absorb shifts in discount rates, the business cycle, and financing conditions that move realizations independently of CGT (e.g., higher real rates or recessions can depress realizations even at a fixed tax rate). I exclude variables such as share repurchases or stock market turnover because they are at least partly post-treatment mediators: a lower CGT can induce substitution toward or away from buybacks and spur trading, so conditioning on them would absorb channels through which the tax operates and would push the estimated elasticity toward a direct (net-of-substitution) effect rather than the policy-relevant total effect.

Table 2 reports the results for $h = 1$ using both top statutory and average effective tax rates.[33] A clear term structure emerges: elasticities are small at one year, peak around the five-year window, and attenuate by the decade horizon. Using the statutory rate, the no-controls elasticities move from -0.11 at $n = 1$ to -0.56 at $n = 5$, easing to -0.38 at $n = 11$; using the effective rate, the elasticity is uniformly larger in magnitude—-0.32, -0.75, and -0.52, respectively. Adding the minimal controls (P_t/D_t and the net-of-dividend tax τ^D) barely changes these patterns (with the exception of the elasticity at $n = 11$ for statutory rates, which becomes as large as for $n = 5$). Bringing in the full macro-finance block reduces magnitudes somewhat but preserves the same hump-shaped profile and the gap between statutory and effective specifications. Standard errors indicate that the medium- and long-run responses are statistically distinguishable from zero and cluster around one-half in absolute value, while the one-year estimates are smaller and noisier, consistent with short-run timing effects that wash out over longer horizons.

The figures are broadly in line with the literature. For 5 years, the numbers are similar to the elasticity used by the Joint Committee on Taxation, which is -0.70 (Gravelle, 2020) and the estimate obtained by Dowd, McClelland, and Muthitacharoen (2015), -0.72. For 10 years, the point estimates are between -0.69 and -0.38 are not entirely dissimilar with the $[-0.5, -0.3]$ band reported by Agersnap and Zidar (2021). This lower elasticity might be related to a wider time window that captures medium-term realizations that were only locked in in the short term (Sarin et al., 2022). While alignment with the literature validates the specification, my primary goal is to obtain an elasticity that can be computed identically in both empirical and model-generated data, providing an additional moment to discipline the estimation.

[33] The tax change is measured over the window $[t-n, t]$, while the realization response is measured over $[t+h-n, t+h]$. Thus, with $h > 0$ the dependent variable intentionally starts after the tax-change window ends. I set $h = 1$ throughout for two reasons. First, with annual data, $h = 0$ estimates are sensitive to the precise dating of reforms and to reform-year timing and aggregation of realizations; setting $h = 1$ is therefore a conservative choice that reduces sensitivity to reform-year measurement. Second, $h = 1$ delivers estimates that are more stable across specifications and more precisely estimated.

Table 2: Empirical macro-estimations of the tax elasticity of realization. *This table reports estimations of β_τ in expression (61) and the implied tax elasticity of realizations (62). It uses aggregate U.S. data on capital gains realizations, the top statutory capital gains tax rate, the effective average marginal capital gains tax rate, and macro-finance controls. The sample period is 1975–2022. For each horizon, τ^* is computed as the average tax rate over the regression sample used for that horizon. Newey–West standard errors are in parentheses.*

	Statutory rate		Effective rate	
	β_τ	$\varepsilon_{\text{tax}}(\tau^*)$	β_τ	$\varepsilon_{\text{tax}}(\tau^*)$
1 year ($n = 1$)	$\tau^* = 0.2243$		$\tau^* = 0.1065$	
	0.39	-0.11	2.70	-0.32
	(0.76)	(0.22)	(1.83)	(0.22)
+ Controls: P_t/D_t, τ^D	-0.44	0.13	2.47	-0.29
	(0.67)	(0.19)	(1.82)	(0.22)
+ Controls: all	0.72	-0.21	3.23	-0.38
	(0.77)	(0.22)	(1.85)	(0.22)
5 years ($n = 5$)	$\tau^* = 0.2140$		$\tau^* = 0.1035$	
	2.07	-0.56	6.51	-0.75
	(1.21)	(0.33)	(2.72)	(0.31)
+ Controls: P_t/D_t, τ^D	2.34	-0.64	6.28	-0.72
	(1.32)	(0.36)	(2.69)	(0.31)
+ Controls: all	2.42	-0.66	5.88	-0.68
	(1.32)	(0.36)	(1.97)	(0.23)
11 years ($n = 11$)	$\tau^* = 0.2119$		$\tau^* = 0.1036$	
	1.40	-0.38	4.50	-0.52
	(1.06)	(0.29)	(2.19)	(0.25)
+ Controls: P_t/D_t, τ^D	2.57	-0.69	5.27	-0.61
	(0.63)	(0.17)	(1.24)	(0.14)
+ Controls: all	1.80	-0.48	4.90	-0.57
	(0.56)	(0.15)	(0.83)	(0.10)

3.3.2. Parameterization

Table 3 contains the parameter values. As anticipated, a subset of parameters is chosen outside of the SMM procedure. In particular, the risk aversion parameter is set to 2, which is a standard value in the literature (e.g., Campbell and Cochrane, 1999, Adam, Marcet, and Beutel, 2017) and ensures full comparability of the results. Moreover, the persistence of the quarterly wage-dividend is calibrated to 0.99, which matches the empirical value of the data for the sample period. To discipline the model further, I calibrated to shocks for dividends, capital taxes, stock buybacks, and the discount factor to perfectly match the

corresponding data counterparts.[34] See Appendix A for details about these series.

The remaining parameters are: the average wage-dividend ratio w, the standard deviation of wage-dividend shocks σ_w, and the parameters characterizing the learning process. To avoid the curse of dimensionality, Kalman gains are obtained as random draws from a log Normal distribution:

$$g^i \sim \log\mathcal{N}\left(\ln(\mu_g) - \frac{1}{2}\ln\left(1+\frac{\sigma_g^2}{\mu_g^2}\right), \ln\left(1+\frac{\sigma_g^2}{\mu_g^2}\right)\right) \tag{63}$$

This formulation allows for heterogeneous gains while dramatically reducing the number of free parameters. There are two natural high-dimensional alternatives. First, one could estimate the Kalman gain agent-by-agent, which requires I free parameters $\{g^i\}_{i=1}^I$. Second, one could estimate the two subjective-noise variances underlying each agent's gain, which requires $2I$ free parameters (e.g., $\{\sigma_{\eta,i}^2, \sigma_{\varepsilon,i}^2\}_{i=1}^I$). Relative to the $2I$-parameter variance benchmark, the approach I follow reduces the number of free parameters from $2I$ to 2 (from 200 to 2 when $I = 100$); relative to the agent-by-agent gain benchmark, it reduces the number from I to 2. Altogether, the parameter vector to be estimated by SMM is $\theta = \{w, \sigma_w, \mu_g, \sigma_g\}$.

Following the extension of SMM by Adam, Marcet, and Nicolini (2016) to include statistics rather than pure moments, the estimated vector $\hat{\theta}$ minimizes the distance between the model $\widetilde{\mathcal{S}}(\theta)$and the empirical data statistics $\hat{\mathcal{S}}$, that is,

$$\hat{\theta} = \underset{\theta\in\Theta}{\arg\min}\left[\hat{\mathcal{S}} - \widetilde{\mathcal{S}}(\theta)\right]'\hat{W}\left[\hat{\mathcal{S}} - \widetilde{\mathcal{S}}(\theta)\right] \tag{64}$$

where $\hat{\mathcal{S}}$ and $\widetilde{\mathcal{S}}(\theta)$ are $M \times 1$ vectors, $\hat{W}$ is a $M \times M$ weighting matrix and Θ is a compact subset of $\mathbb{R}^k$ ($k = 4$ in this case) that contains the true parameter vector θ_0 in its interior. I use a diagonal weighting matrix, with the diagonal entries consisting of the inverse of the individually Newey-West variances of the corresponding observed statistics $\hat{\mathcal{S}}$. This criterion minimizes the sum of squared t-statistics across moments, therefore focusing more directly on matching them. To compute $\widetilde{\mathcal{S}}(\theta)$, I follow a Monte Carlo procedure.

The bottom panel of table 3 contains the estimated $\hat{\theta}$. A few comments are pertinent. First, the estimated values for the gain distribution parameters are $(\hat{\mu}_g, \hat{\sigma}_g) = (0.0277, 0.0057)$, broadly in line with the literature on learning about prices (e.g., Adam et al. (2015), Adam, Marcet, and Nicolini (2016), Adam, Marcet, and Beutel (2017)). As an additional check, appendix Appendix G estimates Kalman gains for synthetic cohorts in the UBS Gallup Survey and fit a log-Normal distribution, finding $(\mu_g^{UBS}, (\sigma_g)^{UBS}) =$

[34] Although equation (37) specifies a highly persistent law of motion for tax rates under agents' subjective measures, the baseline quantitative exercises are conducted over the finite historical sample using the observed tax sequence (with Monte Carlo repetitions over wage shocks). Hence, the targeted moments are finite-sample moments conditional on the realized tax history, rather than population moments under a stationary distribution implied by (37).

Table 3: Parameter values. *The table reports the value of the parameters for the baseline version of the model.*

Non included in SMM			
Persistence wage-dividend ratio	p	0.99	Calibrated
Dividends shocks	ε_t^D		Calibrated
Supply shocks	ε^S		Calibrated
τ^{CG} shocks	$\varepsilon^{\tau^{CG}}$		Calibrated
τ^D shocks	ε^{τ^D}		Calibrated
δ shocks	ε^{δ}		Calibrated
CRRA parameter	γ	2	Literature
Included in SMM			
Average wage-dividend ratio	w	6.02	Estimated
Standard deviation WD shock	σ_W	0.0203	Estimated
Kalman gain mean	μ_g	0.0277	Estimated
Kalman gain dispersion	σ_g	0.0057	Estimated

$(0.0293, 0.0062)$. Thus, the SMM estimates look remarkably close to these points estimates from surveys, Additionally, a key point of models with learning about prices is that they generate enough price volatility without resorting to either large risk aversion or large fundamental risk. In this vein, the estimated standard deviation of the wage shocks is 0.0203, which is close to the empirical point estimate (in the sample, a point estimate for σ_w is 0.0215, but 0.0203 is within the confidence bands).

Finally, the combination of PEA for solving the model with SMM for parameter estimation requires careful implementation. PEA proceeds by fixing the structural parameters θ and iterating on the approximation until convergence, whereas SMM evaluates many candidate values of θ, each requiring a new set of PEA parameters. Thus, PEA constitutes an inner fixed-point routine nested within SMM's outer minimization over θ. Appendix E describes an efficient algorithm for managing this PEA-SMM loop.

3.4. Model evaluation

This section assesses the model's ability to replicate a comprehensive set of empirical facts. I distinguish between moments used in estimation and a separate set of "validation moments" that serve as out-of-estimation checks. This approach mirrors the distinction between training and validation sets in machine learning and provides a more stringent test of the model's structural validity. If the model captures the underlying economic mechanisms—rather than merely fitting a particular set of moments—it should also match statistics that were not directly targeted in estimation.

Table 4 presents the results. The top panel reports the eight statistics included in the SMM estimation. The model matches these moments well: t-statistics for the null hypothesis that model and data moments are equal are uniformly small in absolute value, with none exceeding conventional significance thresholds. The model successfully captures the key "excess volatility" facts that have long puzzled asset pricing researchers: high volatility of the price-dividend ratio (19.90 vs. 19.12 in the data), the Campbell-Shiller decomposition showing that most valuation movements reflect expected returns rather than cash flows, and the high persistence of valuations. Crucially, the model also matches the 11-year tax elasticity of realizations (-61.52 vs. -60.07), ensuring that investors' behavioral responses to taxation align with the empirical evidence. This is essential for the credibility of the counterfactual tax experiments that follow.

The bottom panel reports the validation moments—statistics that were deliberately excluded from the estimation criterion. The model performs remarkably well on this validation set. First, it generates the correct negative correlation between valuations and dividend tax rates (-75.94 vs. -67.19), capturing the upward trend in valuation that coincides with the downward trend in capital taxes. Second, the model replicates the one-year impulse response of valuations to price shocks (2.70 vs. 3.21), demonstrating that the learning dynamics generate realistic momentum in the price-dividend ratio at business-cycle frequencies, without generating overreaction to shocks. Third, the 5-year tax elasticity (-78.12 vs. -72.45)—which captures medium-run realization responses distinct from the long-run elasticity used in estimation—falls within the confidence interval of empirical estimates, suggesting the model correctly captures the term structure of behavioral responses to taxation. Fourth, the model is able to reproduce the positive association between trading and price changes, providing qualitative validation of the model's trading mechanisms even if the quantitative magnitude is somewhat overstated. Finally, dividend growth and price growth are included as internal consistency checks; mean and standard deviation of dividend growth disciplined by calibration, and mean price growth largely implied by the fit to mean returns and valuations.

4. The effects of tax cuts

In this section, the effect of tax cuts on the stock market is evaluated. Leveraging the good quantitative performance of the model, I use it to run experiments that quantify the effect of tax cuts on different statistics. In Section 4.1., I outline the causal effect measure and decompose it. Section 4.2.explores the sensitivity of the relative average effect of tax cuts to key parameters and model assumptions. Finally, in Section 4.3., I present a calibration exercise.

Table 4: Quantitative performance of the model. *This table reports the M statistics included in the SMM estimation. The first two columns report the observed statistics for the US data. The next four columns report model-implied statistics and their t-statistics using table 3's parameters. Rates of growth and returns are annualized. Variances, covariances, correlations and elasticities have been multiplied by 100. The tax elasticities of realizations corresponds to the values for the specifications with the effective tax rate and* $(P_t/D_t,\ \tau_t^D)$ *as controls in Table 2.*

	US data	Model	
	$\hat{S}_j$	$\widetilde{S}_j(\hat{\theta})$	t-stat
	Included in the SMM estimation		
$\mathbb{E}(PD_t)$	44.73	47.08	-0.63
$\mathbb{E}(r_t^s)$	8.97	8.81	0.07
$\text{Var}(\ln PD_t)$	19.12	19.90	-0.22
$\text{Cov}(\ln PD_t, \bar{r}_t)$	-14.52	-15.14	0.22
$\text{Cov}(\ln PD_t, \bar{d}_t)$	4.60	4.75	-0.18
$\sigma(\ln R_t^s)$	15.00	15.62	-0.45
$corr(PD_t, PD_{t-1})$	98.12	96.23	0.43
$\varepsilon_{\text{tax},11y}(\tau^*)$	-60.07	-61.52	0.08
	Non-included in the SMM estimation		
$corr(PD_t, \tau_t^D)$	-67.19	-75.94	0.88
$\mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 1, I_t) - \mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 0, I_t)$	3.21	2.70	1.33
$\varepsilon_{\text{tax},5y}(\tau^*)$	-72.45	- 78.12	0.18
$\text{corr}(TV_t, \lvert P_t/P_{t-1} - 1\rvert)$	32.30	51.24	-1.50
$\mathbb{E}(P_t/P_{t-1} - 1)$	6.13	6.02	0.02
$\mathbb{E}(D_t/D_{t-1} - 1)$	2.50	2.50	0.00
$\sigma(\ln \frac{D_t}{D_{t-1}})$	3.40	3.40	0.00

4.1. Measuring the Causal Effect of Tax Cuts

In this section, I introduce a measure of the effect of tax cuts on stock market volatility. Leaning on the good quantitative properties of the model, I use counterfactual simulations of scenarios without tax cuts.

Let $\tilde{S}_j(\hat{\theta} \mid \{\tau_t\})$ denote a model-implied statistic j computed under the empirical tax path $\{\tau_t\}_{t=1}^T$ at the estimated parameter vector $\hat{\theta}$. I measure the causal effect of the observed tax changes by comparing this statistic to a counterfactual economy in which taxes are held fixed at a benchmark level $\bar{\tau}$ for all t.[35] Formally, define the tax-path effect

$$\Delta S_j \;\equiv\; \tilde{S}_j(\hat{\theta} \mid \{\tau_t\}) - \tilde{S}_j(\hat{\theta} \mid \{\bar{\tau}\}) \tag{65}$$

A positive ΔS_j indicates that, in the model, the observed sequence of taxes raises statistic j relative to the constant-tax benchmark $\bar{\tau}$.

To express magnitudes relative to the data, let $\hat{S}_j$ denote the empirical counterpart of statistic j (e.g., for $j = \mathrm{Var}(\ln PD_t)$, $\hat{S}_j$ is the sample variance of the price–dividend ratio). I report the dimensionless relative effect

$$\frac{\Delta S_j}{\hat{S}_j} \tag{66}$$

which, when multiplied by 100, gives the percent of the observed level of statistic j accounted for by moving from the constant benchmark $\bar{\tau}$ to the realized tax path $\{\tau_t\}$. This object summarizes the effect of the entire sequence of tax changes relative to $\bar{\tau}$.

To summarize the magnitude of the tax reform itself, define the average absolute deviation of the tax path from the benchmark,

$$\bar{g}_\tau \;\equiv\; \frac{1}{T}\sum_{t=1}^{T} \left|\tau_t - \bar{\tau}\right| \tag{67}$$

If $\bar{\tau}$ is chosen as the maximal tax rate in the sample, then $\bar{g}_\tau$ corresponds to the typical cut size (in percentage points) relative to that initial level. Finally, normalizing the relative volatility effect by $\bar{g}_\tau$ yields a semi-elasticity per one percentage point:

$$\varepsilon_j^\tau \;\equiv\; \frac{\Delta S_j}{\hat{S}_j} \cdot \frac{1}{\bar{g}_\tau} \tag{68}$$

This semi-elasticity maps a one-percentage-point average tax deviation from $\bar{\tau}$ into a proportional change in statistic j. For example, if $\bar{\tau}$ is the maximal tax rate, $\varepsilon_j^\tau = 0.30$ means

[35] A natural concern is that the historical tax sequence combines shifts in tax levels with time-variation in taxes (policy innovations), which are themselves a direct source of volatility in the model. Section 4.2 isolates the level (trend) component by shutting down tax innovations while holding fixed the remaining shock history; the resulting volatility effects are similar.

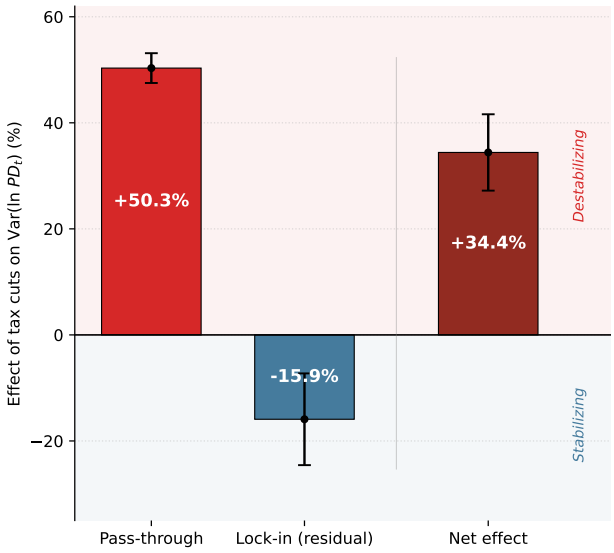


Figure 3: Causal effect of tax changes in volatility. *The first column shows the relative causal effect coming from the pass-through effect; the second column shows the residual effect, that includes the lock-in effect and general equilibrium interactions. The third column shows the net effect, which is the sum of the previous two as in equation (74). The counterfactual tax rate is* $\bar{\tau} = \tau_{1975} = 16.32\%$.

that a 1pp average tax cut (relative to $\bar{\tau}$) is associated with a 0.30% increase in the statistic j measured relative to its empirical level $\hat{S}_j$.

In the exercise, $\bar{\tau}$ is kept at its 1975 value, which is representative of the "high" tax regime and focus on the variance of the log price-dividend ratio, which I refer to as $\tilde{S}_{\text{Var}}(\hat{\theta} \mid \{\bar{\tau}\})$. Under this counterfactual, the model predicts volatility of $\tilde{S}_{\text{Var}}(\hat{\theta} \mid \{\bar{\tau}\}) = 13.32$, compared to $\tilde{S}_{\text{Var}}(\hat{\theta} \mid \{\tau_t\}) = 19.90$ under the historical path. The relative causal effect is therefore

$$\frac{\Delta S_{\text{Var}}}{\hat{S}_{\text{Var}}} = \frac{19.90 - 13.32}{19.12} = 0.3441$$

That is, absent the tax cuts, volatility would have been approximately 34% lower on average, with a 95% confidence interval ranging from 27% to 41%.[36] The third bar of figure 3 plots it. The implied semi-elasticity is $\varepsilon^{\tau}_{\text{Var}} = 6.17$: a one percentage point tax cut raises volatility by roughly 6% relative to the empirical benchmark. Taken together, these results suggest that the cumulative effect of capital gains tax cuts over this period represents a first-order contributor to stock market volatility.

[36] Alternatively, one could normalize ΔS_j by the model-implied statistic rather than its empirical counterpart. In this case, the denominator becomes 19.90, yielding a relative effect of 33.06%. The modest discrepancy between the two normalizations reflects the small gap between the empirical and model-based point estimates.

4.1.1. Decomposing the Causal Effect

In this section, I decompose the volatility response to the historical tax cuts into (i) a pass-through benchmark and (ii) a residual that captures the incremental contribution of realization-based taxation. The pass-through benchmark is disciplined by a counterfactual economy in which capital gains are taxed on an accrual basis under the same tax path $\{\tau_t\}$. Relative to the baseline realization system, the accrual base eliminates trading-contingent realization incentives—shutting down the kink and the associated inaction region—while preserving the scaling of after-tax payoffs implied by the statutory rate. The difference in volatility responses between the historical tax path $\{\tau_t\}$ and a high constant-rate benchmark $\{\bar{\tau}\}$, computed within the accrual economy, therefore provides a transparent measure of how taxes affect volatility through capitalization into prices in the absence of lock-in.

Isolating "lock-in" as a separate, additively identified mechanism is less straightforward in the quantitative model. Unlike in the analytic model—where one can remove pass-through by holding expectations fixed conditional on the discrete state—here beliefs evolve endogenously over a continuum of states through learning and heterogeneous updating. Importantly, the belief updating equations are identical across the realization and accrual economies; what differs is only the tax system (realization-based versus accrual-based) and the equilibrium objects it disciplines.[37] Because trading affects prices and prices feed back into beliefs, changing realization incentives necessarily changes the equilibrium price signal and hence the endogenous path of beliefs. For this reason, I interpret the difference between the total effect in the baseline realization economy and the pass-through benchmark measured in the accrual economy as a residual: it captures the incremental effect of implementing the same statutory tax changes under a realization base rather than an accrual base, inclusive of general-equilibrium interactions between trading, prices, and belief dynamics. The sensitivity analysis in Section 4.2 complements this accounting by attenuating learning to verify that the pass-through component behaves in the predicted direction.[38]

Accrual-based taxation. Consider an asymmetric accrual regime, where gains are taxed and losses are not subsidized but are carried forward to offset future gains. In

[37] One might worry that policy-induced changes in volatility would alter agents' learning speeds (a Lucas critique). In constant-gain learning, however, the gain parameters are structural primitives of the updating rule and are therefore taken to be policy-invariant.

[38] Throughout this decomposition, all comparisons are conducted along a common realized history of exogenous shocks and primitives: the dividend/fundamental innovations and the net supply process $\{\tilde{S}_t\}$ are held fixed across tax paths and across the baseline (realization-based) and accrual counterfactuals. The belief-updating equations and information structure are also held fixed. What differs across legs is solely the tax treatment of gains (realization-based vs. accrual) and the associated equilibrium objects it disciplines—prices, portfolios, realized gains, and tax-base states (e.g., basis and loss carryovers) adjust endogenously within each economy given the same shock realization.

this case, the budget constraint of the investor reads as

$$C_t^i + P_t S_t^i = W_t + ((1-\tau_t^D)D_t + P_t)S_{t-1}^i - \tau_t^A \max\{P_t - P_{t-1} - L_t^i, 0\}S_{t-1}^i + T_t^i \quad (69)$$

The stock of unrealized losses is a state variable whose law of motion is:

$$L_{t+1}^i = \frac{S_{t-1}^i}{S_t^i} \max\{L_t^i - (P_t - P_{t-1}), 0\} \quad (70)$$

With accrual taxation, the problem is differentiable everywhere, which eases computations. Analogously to the realization-based economy, let $\psi^A(z_{A,t}^i)$ summarize investor i's discounted subjective marginal value of carrying one additional share into $t+1$ relative to the current price where $z_{A,t}^i$ is the state vector under accrual taxation. Then, the KKT conditions boil down to

$$\begin{aligned} C_t^i &\geq (\psi_t^A)^{-1/\gamma} \text{and } \; S_t^i = \bar{S} \\ C_t^i &= (\psi_t^A)^{-1/\gamma} \text{and } \; S_t^i \in (0, \bar{S}) \\ C_t^i &\leq (\psi_t^A)^{-1/\gamma} \text{and } \; S_t^i = 0 \end{aligned} \quad (71)$$

where $\psi_t^A \equiv \psi^A(z_t^i)$. The model can be solved by a simplified version of the algorithms used to solve the realization-based economy. See Appendix E for details.

Channels decomposition. Let the causal effect in the accrual-based economy be

$$\Delta S_j^A \;\equiv\; \tilde{S}_j^A(\hat{\theta} \mid \{\tau_t\}) - \tilde{S}_j^A(\hat{\theta} \mid \{\bar{\tau}\}) \quad (72)$$

where $\tilde{S}_j^A$ refers to the moment generated in the accrual-based economy. Because accrual taxation is not trade-contingent, it eliminates the kink in the household problem and the associated no-trade region induced by realization incentives. The effect ΔS_j^A therefore provides a transparent benchmark for how taxes affect volatility through capitalization into equilibrium prices when realization-driven trading distortions are absent.[39]

To relate this benchmark to the total effect in the baseline realization economy, define the residual difference

$$\begin{aligned} \Delta S_j^R \;\equiv\; \Delta S_j - \Delta S_j^A = &\Big[\tilde{S}_j\big(\hat{\theta} \mid \{\tau_t\}\big) - \tilde{S}_j^A\big(\hat{\theta} \mid \{\tau_t\}\big)\Big] \\ &- \Big[\tilde{S}_j\big(\hat{\theta} \mid \{\bar{\tau}\}\big) - \tilde{S}_j^A\big(\hat{\theta} \mid \{\bar{\tau}\}\big)\Big] \end{aligned} \quad (73)$$

Equation (73) compares the response of the statistic to tax cuts under realization-based taxation to the corresponding response in the accrual counterfactual. The first bracket

[39] In the accrual pass-through benchmark, I apply the same capital-gains tax path $\{\tau_t\}$ from the baseline realization economy to the accrual tax base (i.e., $\tilde{S}_j^A(\hat{\theta} \mid \{\tau_t\})$ uses the identical $\{\tau_t\}$ sequence). See Section 5.2. for a revenue-neutral approach to accrual-based taxation.

measures the gap between realization and accrual outcomes along the observed tax path, i.e., the quantitative importance of realization incentives in the estimated economy. The second bracket subtracts the analogous gap evaluated at the benchmark tax level. The resulting difference, ΔS_j^R should be interpreted as a residual component that captures the contribution of the lock-in effect along with general-equilibrium interactions with belief formation and price dynamics. In particular, because beliefs govern both valuation and trading motives in the quantitative model, changes in equilibrium prices and belief dynamics can alter when the realization kink binds; accordingly, ΔS_j^R is not a standalone estimate of a pure lock-in channel.

With this notation, the total causal effect admits the accounting identity

$$\Delta S_j = \Delta S_j^A + \Delta S_j^R \tag{74}$$

I refer to ΔS_j^A as the pass-through benchmark (the effect of taxes absent realization-based lock-in), and to ΔS_j^R as the residual contribution associated with realization incentives and interactions.

Columns 1 and 2 of Figure 3 report these two objects. In the accrual-based economy, lower taxes are highly destabilizing, generating an average increase in volatility of 50%, with a 95% confidence interval of [47.50%, 53.13%]. In the baseline realization economy, the corresponding total effect is smaller because realization incentives partially offset the destabilizing accrual benchmark. Quantitatively, the residual component ΔS_j^R is approximately -15%, with a confidence interval of $[-24.56\%, -7.26\%]$, indicating that realization-based trading distortions—together with their equilibrium interactions—attenuate, but do not overturn, the destabilizing pass-through benchmark.

4.1.2. Comparison to Dai-Shackelford–Zhang (2013)

While the paper's main focus is on volatility of prices/valuations, dai2013capital study the volatility of returns. In this section I apply the same toolkit to returns and compare the model's predictions to their estimates.

Empirical benchmark (Dai–Shackelford–Zhang). In their baseline DiD regressions, they find increases of about 14% (RA–1978) and 18% (TRA–1997) in the monthly standard deviation for large–gain portfolios relative to small–gain portfolios; and 9% (RA-1978) and 17% (TRA - 1997) for non–dividend payers relative to dividend payers. Converting these to semi–elasticities per 1 percentage–point (pp) tax change by dividing by the statutory cut magnitudes (RA: 35%$\to$28%, $\approx$ 7 pp; TRA: 28%$\to$20%, $\approx$ 8 pp) yields approximately +2.0–2.3% per 1 pp for the large–gain split and +1.3–2.1% per 1 pp for the non–dividend split. These are treated–minus–control effects measured over the pre/post windows defined

around each reform.

Model-based aggregate effect. Let $\tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\tau_t\})$ denote the model-implied standard deviation of stock returns under the historical tax path and $\tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\bar{\tau}\})$ the counterpart under a constant benchmark $\bar{\tau}$. Keeping $\bar{\tau}$ at its 1975 value, the model yields $\tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\bar{\tau}\}) = 11.20$ versus $\tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\tau_t\}) = 15.62$ under the historical path. The relative causal effect is

$$\frac{\tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\tau_t\}) - \tilde{S}_{\sigma_R}(\hat{\theta} \mid \{\bar{\tau}\})}{\hat{S}_{\sigma_R}} = \frac{15.62 - 11.20}{15.00} = 0.2947$$

i.e., absent the tax cuts, volatility would have been about 30% lower on average (95% CI: [28.21%, 30.74%]). The implied semi-elasticity is $\varepsilon^{\tau}_{\sigma_R} = 5.29$: a one–percentage–point tax cut raises the standard deviation of returns by roughly 5% relative to the empirical benchmark. Despite the very different estimands (aggregate path effect vs. cross-sectional DiD), the magnitudes are of the same order.

Local, reform-specific effects in the model. Because Z. Dai, Shackelford, and Zhang (2013) analyze discrete reforms, I also compute model-based local semi-elasticities using their time windows: for RA–1978 (January 1976–May 1981) I obtain $\varepsilon^{\tau}_{\sigma_R} = 5.49$, and for TRA–1997 (January 1994–December 2000) $\varepsilon^{\tau}_{\sigma_R} = 2.29$. Thus, the model predicts a larger effect for the 1978 cut and an elasticity closely aligned with Z. Dai, Shackelford, and Zhang (2013) for the 1997 reform. This broad alignment supports the credibility of the mechanism and suggests the model might be useful for interpreting their cross-sectional evidence.

4.1.3. Other statistics

The previous sections have focused on the effect of tax cuts on various measures of market volatility. This section complements the analysis by examining the effects on other statistics. The absence of CGT cuts would have reduced mean stock valuations by approximately 20%. Similarly, the correlation between valuations and dividend taxes is diminished by about 25%, as the absence of CGT cuts softens the upward trend in valuations that drives this correlation. Thus, tax cuts contributed to stock market exuberance, although the effect on valuation levels is considerably smaller than the effect on volatility.

The increase in volatility can be decomposed into the two Campbell–Shiller components. The effect is much stronger for the comovement between valuations and future returns than for the comovement between valuations and future dividend growth, suggesting that factors other than fundamental growth played an important role in explaining the heightened volatility. This finding is entirely consistent with the model: lower taxes accentuate the price-expectations feedback loop, which translates into volatility unexplained by fundamentals—captured by the covariance of valuations with future returns. Consistent

with this interpretation, the one-year-ahead IRF falls from 2.70 to 1.87, which corresponds to a fall of about 25% with respect to the empirical benchmark 3.21. In turn, higher taxes would have slightly increased the serial correlation of valuations, partially offsetting the stabilizing effect.

Finally, the model suggests that tax cuts played a central role in generating the co-movement between trading volume and price changes. Higher taxes would have reduced this correlation by half when ΔS_j is normalized by the data moment, and by 32% when normalized by the model moment $\tilde{S}_j(\hat{\theta} \mid \{\tau_t\})$.[40] This finding is consistent with the view that CGT reduces trading: an equivalent change in prices would be associated with lower trading volume under a high-tax regime.

Table 5: Causal Effect of Capital Gains Tax Cuts on selected statistics. *The first column report the moments. The second and third columns report the model-implied statistics using the parameterization in Table 3 and using the empirical tax sequence and the constant tax rate $\bar{\tau}$ respectively. The last column report the relative causal effect as defined in equation (66).*

	$\tilde{S}_j(\hat{\theta} \mid \{\tau_t\})$	$\tilde{S}_j(\hat{\theta} \mid \{\bar{\tau}\})$	$\Delta S_j / \hat{S}_j$ (%)
$\mathbb{E}(PD_t)$	47.08	37.64	21.09
$\text{Var}(\ln PD_t)$	19.90	13.32	34.41
$\text{Cov}(\ln PD_t, \bar{r}_t)$	-15.14	-9.53	38.64
$\text{Cov}(\ln PD_t, \bar{d}_t)$	4.75	3.78	21.08
$\text{corr}(PD_t, PD_{t-1})$	96.23	97.95	-1.75
$\text{corr}(PD_t, \tau_t^D)$	-75.94	-59.55	24.40
$\mathbb{E}(PD_{t+4} \mid \varepsilon_t^P = 1, I_t) - \mathbb{E}(PD_{t+4} \mid \varepsilon_t^P = 0, I_t)$	2.70	1.87	25.80
$\text{corr}(TV_t, \lvert P_t/P_{t-1} - 1 \rvert)$	51.24	34.65	51.37
$\mathbb{E}(P_t/P_{t-1} - 1)$	6.02	5.37	10.56

4.2. Sensitivity Analysis

In this section, I explore the sensitivity of the relative causal effect to alternative modeling choices. Each exercise uses the baseline parameterization outlined in Table 3, isolating the marginal contribution of the assumption under consideration. Table 6 collects the results.

Tax level versus tax volatility. The baseline comparison evaluates the historical tax path against a counterfactual with high and constant taxes. Because the estimated tax process is highly persistent, this counterfactual removes not only the wedge implied by a lower average tax rate but also the stochastic innovations to taxes (i.e., tax policy risk). Therefore, the measured causal effect could in principle partly reflect the elimination of transitory tax fluctuations (for example, the sharp cut in the early 1980s that was partially

[40] These different normalizations are immaterial for most moments, as the model matches the data moments closely. For this particular correlation, the difference is material because the model overstates the correlation.

reversed within a few years) rather than the effect of a sustained shift in the tax level, which is the object of interest.

To separate these components, I decompose the observed tax series into a long-run trend, $\{\tau_t^{\text{trend}}\}$, capturing the gradual transition from high to low taxes, and a residual capturing short-run fluctuations around that trend. I extract $\{\tau_t^{\text{trend}}\}$ using an HP filter with a high smoothing parameter so that the trend abstracts from transitory policy reversals. This yields the following accounting decomposition of the total effect into a "trend (level) effect" and a "volatility (risk) effect":

$$\big(\tilde{S}_j(\hat{\theta}) \mid \{\tau_t\}\big) - \big(\tilde{S}_j(\hat{\theta}) \mid \{\bar{\tau}\}\big) = \underbrace{\big(\tilde{S}_j(\hat{\theta}) \mid \{\tau_t^{\text{trend}}\}\big) - \big(\tilde{S}_j(\hat{\theta}) \mid \{\bar{\tau}\}\big)}_{\text{trend effect}} + \underbrace{\big(\tilde{S}_j(\hat{\theta}) \mid \{\tau_t\}\big) - \big(\tilde{S}_j(\hat{\theta}) \mid \{\tau_t^{\text{trend}}\}\big)}_{\text{volatility effect}} \tag{75}$$

The first term measures how volatility would change if the tax regime followed only the secular decline from its 1975 anchor, while the second captures the additional volatility generated by short-run tax movements around that trend.

Table 6 reports results for $\lambda_{\text{HP}} = 1{,}000{,}000$ and $\lambda_{\text{HP}} = 100{,}000$. Despite the aggressive smoothing of the tax path, the relative causal effect is robust: for the lower smoothing parameter, it remains virtually unchanged at 34.38%; for the higher smoothing parameter, it declines modestly to 30.16%. These results suggest that short-run tax fluctuations contribute only a limited share of the total effect; most of the volatility impact is accounted for by the persistent decline in tax levels. Hence, the causal estimate is not primarily driven by tax-policy-risk innovations, but by the sustained shift in the tax regime, in line with the mechanism emphasized in Section .

No dividend tax cuts. What would the effect of capital gains tax cuts be in the absence of dividend tax cuts? To answer this question, I hold $\tau_t^D = \bar{\tau}^D \equiv \max\{\tau_t^D\}$ for all t and recompute the causal effect of capital gains tax cuts. In this counterfactual, $\tilde{S}_{\text{Var}}(\hat{\theta} \mid \{\tau_t^{CG}, \bar{\tau}^D\}) = 17.35$, so keeping dividend tax rates high throughout has mild effects on volatility. However, keeping both tax rates high and constant delivers very low volatility: $\tilde{S}_{\text{Var}}(\hat{\theta} \mid \{\bar{\tau}, \bar{\tau}^D\}) = 7.54$. Thus, the effect of capital gains tax cuts is larger in this scenario, delivering a relative effect of approximately 50%. This result suggests a remarkable interaction between capital gains and dividend taxes: keeping both high would have reduced valuation volatility substantially.

Statutory rates. The baseline uses average marginal effective capital-gains tax rates. While standard in the literature as a proxy for the aggregate tax burden, effective rates are not a pure policy object: they also embed time variation in market composition (taxable versus tax-exempt ownership), which may itself respond endogenously to tax reforms. To isolate policy variation more cleanly, I therefore re-compute the counterfactuals using the top statutory capital-gains tax rate. Furthermore, statutory rates avoid imposing a population-average effective wedge uniformly across heterogeneous investors; all taxable

investors face the same statutory schedule, so cross-agent heterogeneity is not mechanically collapsed into a single time-varying effective rate. Using statutory rates does not substantially affect the size of the causal effect; the point estimate decreases slightly to 31.46%. The more sizable change concerns the semi-elasticity: since statutory tax rates display larger variation, the average tax cut is larger and the elasticity is correspondingly smaller, averaging 2.55.

Alternative assumptions about taxes. The tax rates used in the baseline economy reflect several modeling choices; here I explore the sensitivity to some of them. First, I introduce tax rebates, whereby agents receive a refund of the taxes they paid in the previous period.[41] Rebates play a paramount role in the risk-sharing view of capital gains taxation, making this a natural test. According to this view, lower taxes reduce risk-sharing, thereby increasing volatility; thus, a larger effect of tax cuts is expected. In the model, including tax rebates is relatively inconsequential: the causal effect declines modestly to 30.29%. Moreover, the model does not generate a larger effect, so the risk-sharing hypothesis prediction does not find much support—perhaps unsurprisingly, given the asymmetric and realization-based nature of the tax system.

Second, the baseline economy assumes that agents perceive tax rates as unit root processes. I depart from this assumption in two ways. I begin by exploring the possibility of tax foresight, as evidence suggests that markets have substantial ability to forecast taxes (Kueng (2014)). To address this possibility, I endow investors with perfect foresight of near-term taxes (which are indeed the only relevant ones under learning about prices), such that $\mathbb{E}_t^{\mathcal{P}}(\tau_{t+1}^{CG}) = \tau_{t+1}^{CG}$ and $\mathbb{E}_t^{\mathcal{P}}(\tau_{t+1}^{D}) = \tau_{t+1}^{D}$. It turns out that the effect of near-term foresight is very modest, delivering a relative effect just slightly below the baseline: 31.46% versus 34.41%. Alternatively, I let agents learn about taxes using the same constant-gain scheme they employ to learn about prices and dividends. Learning about taxes makes expected taxes somewhat smoother but also heterogeneous, given the different gains among investors. Overall, this additional disagreement reduces the relative effect of taxes, as greater disagreement induces more trading and thereby reinforces the lock-in effect.

Discount factor and supply shocks. The quantitative model incorporates calibrated series for the discount factor and aggregate stock supply to match the secular decline in risk-free rates and the rise in stock repurchases—factors that may independently contribute to higher volatility[42] As a robustness check, I compute the relative effect of tax cuts under a constant discount factor equal to the sample mean of the calibrated series (0.9986). In this case, the effect of tax cuts diminishes to 29.38%. That is, tax cuts would have been somewhat less consequential for volatility in a world with higher and

[41] This is a convenient simplification that preserves recursivity in the system; otherwise, tax rebates and prices are simultaneously determined, which represents a deeper deviation from the no-rebates case and potentially compromises comparability.

[42] See, for instance, Adam (2020) about the decline in safe rates.

stable interest rates, although the attenuation is modest.[43] When aggregate supply is held constant at unity throughout (eliminating stock repurchases), the effect of tax cuts falls further to 22.58%. The model thus suggests that stock repurchases interacted significantly with capital gains tax cuts to amplify stock market volatility.

Price-expectations feedback loop. A central ingredient of the model is learning about prices, which is the source of the belief–price feedback underlying the pass-through channel. To isolate the quantitative role of this feedback from the effects of the realization-based tax wedge, I conduct counterfactual simulations that keep the tax schedule and all other primitives (including the Kalman gains for dividend growth) fixed while varying only the Kalman gain in the capital gains belief updating equation. Specifically, I set the mean gain at 0.001, which makes capital gains beliefs almost unresponsive to price signals and therefore largely shuts down the feedback loop.

In this case, the relative causal effect of the historical tax cuts on volatility falls to 4.62%. Recomputing the lock-in vs. pass-through decomposition under the same low-gain calibration shows that the pass-through component declines from 50% in the baseline to 31.14%, consistent with a substantially weaker amplification of belief fluctuations into prices. At the same time, the lock-in residual becomes more negative, moving from -15% to -26.52%, so the two components nearly offset each other. This pattern suggests that the pass-through critically depends on the strength of price-belief feedback, and it also highlights that the magnitude of the residual depends on general-equilibrium interactions with that feedback.

Beliefs heterogeneity. In addition, I reduce heterogeneity in learning gains to $\sigma_g = 0.001$. This change modestly increases the effect to about 38%, consistent with weaker trading motives and a less stringent lock-in mechanism when cross-sectional belief dispersion is compressed.

4.3. A Historical Counterfactual

I illustrate the effect of tax cuts through a calibration exercise. Instead of using discount factor shocks to replicate the empirical dynamics of the risk-free rate, I calibrate a series of δ shocks to match the empirical dynamics of the price-dividend ratio exactly. This approach ensures that the model replicates the data precisely when fed the empirical time series of taxes, offering a complementary angle to the previous results: it allows for an exploration of the timing of tax cut effects on market valuations.

Figure 4 illustrates the results. The red line represents the estimated model with calibrated δ shocks, which matches the empirical data by construction. The blue line depicts the mean of simulations under a scenario where $\tau_t^{CG} = \tau_{1975}^{CG}$ in every period. Interestingly, until the second half of the 1990s, the series with and without tax cuts

[43] An alternative interpretation of the baseline result is that lower interest rates may have more muted effects on market volatility—and hence financial stability—when capital gains taxes remain relatively high.

Table 6: Sensitivity analysis. *The table reports the relative causal effect of tax cuts and semi-elasticities according to equations (66) and (68) for alternative specifications of the model. It uses* $\bar{\tau} = \tau_{1975}$ *and focuses on the variance of the log price-dividend ratio* $\tilde{S}_{Var}(\hat{\theta} \mid \cdot)$*. Tax trend smooth uses* $\lambda_{HP} = 100{,}000$*; the very smooth one uses* $\lambda_{HP} = 1{,}000{,}000$*. The small average gain sets* $\mu_g = 0.001$ *and the small gain standard deviation uses* $\sigma_g = 0.001$*.*

	Baseline	**Smooth Trend**	**Very Smooth Trend**	**No τ^D cuts**
Relative Effect	34.41	34.38	30.16	51.27
Semi-elasticity	6.17	6.17	5.41	9.20
	Constant δ	**No supply shocks**	**Small μ_g for prices**	**Small σ_g**
Relative Effect	29.38	22.58	4.62	37.81
Semi-elasticity	5.27	4.05	0.83	6.78
	Tax Rebates	**Statutory rates**	**Tax Foresight**	**Tax Learning**
Relative Effect	30.29	31.46	31.46	26.39
Semi-elasticity	5.43	2.55	5.64	4.74

appear remarkably similar, despite several tax cuts having been implemented since the 1980s. This observation can be attributed to nonlinearities in the model: the relationship between taxes and volatility is highly convex, such that tax cuts have particularly large effects when tax rates are already low. Since the mid-1990s, the economy with higher taxes underperforms, opening a valuation gap that only begins to close toward the end of the sample. Thus, in the model, a scenario without tax cuts would have reduced stock market exuberance significantly.

The model without tax cuts implies valuations that are considerably more stable, featuring a much smoother dot-com boom and a far less pronounced decline during the Great Recession. In contrast, the COVID-19 boom-bust episode appears relatively insensitive to tax policy. These differences point toward potentially distinct drivers: the model suggests that a substantial portion of the dot-com boom and the Great Recession crash were driven by extrapolative beliefs, so that higher taxes would have tempered these swings; in contrast, the COVID-19 episode was more fundamentals-driven, so capital gains taxes do comparatively little to stabilize it.

Yet, even with high taxes, the model predicts a boom-and-bust pattern around the 2000s, suggesting that other forces are necessary to fully explain the bubble. For example, Bordalo et al. (2024) proposed that overreaction in long-term expectations was a key element in explaining the bubble. Moreover, in terms of levels, the model without tax cuts still delivers a sustained run-up in valuations from 20 to 50, albeit not reaching the observed peak. This suggests that complementary explanations—such as the reduction in macroeconomic risk boosting demand for risky assets proposed by Lettau, Ludvigson, and Wachter (2008)—may also play a role in accounting for the full extent of the sustained

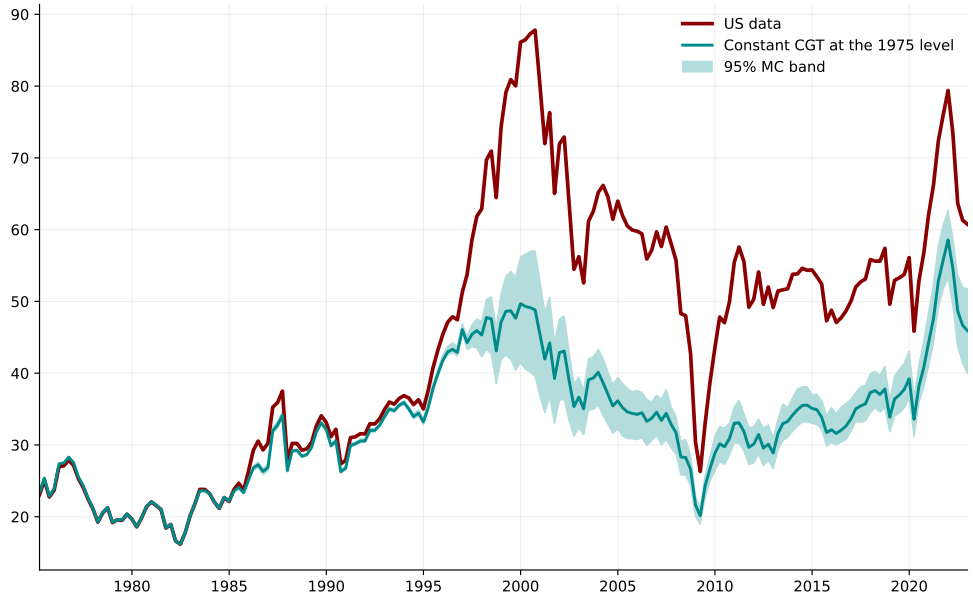


Figure 4: The price-dividend ratio without tax cuts in a calibrated version of the model.

increase in market valuations.

5. Policy experiments with alternative fiscal instruments

In this section, I use the model to evaluate how two alternative tax instruments that have featured prominently in recent policy discussions—an accrual-based capital gains tax (CGT-A) and a financial transactions tax (FTT)—affect stock market volatility. I first map these instruments into the same analytic framework developed in Section 2 to derive predictions for how each policy shapes market volatility. I then quantify these effects in two experiments. The first implements a revenue-neutral reform that replaces realization-based taxation with an accrual-based regime, that compares realization vs accrual regimes. The second introduces, on top of the baseline realization-based CGT, a small FTT and a small CGT-A and compares their performance.

5.1. Analytic insights

In this section, I study the effects of a tax on unrealized capital gains and on financial transactions on volatility using the analytic model in Section 2.

5.1.1. Accrual-based capital gains taxation

In this section, I show that accrual-based taxation is appealing for a simple reason: it retains the stabilizing pass-through effect while avoiding the destabilizing lock-in effect.

To show this, consider the same structure described in Section 2 for a moment. Instead of a tax on realized-gains, I study a tax that applies symmetrically to unrealized gains and losses. Moreover, I focus on a knife-edge analytical case where beliefs satisfy $\pi_A = 1 - \pi_B$. This setup delivers a sharp characterization that cleanly isolates the mechanism at play. In the quantitative experiments, I relax beliefs assumptions and implement asymmetric taxation with carryover losses.

In this case, the budget constraint simplifies to

$$C_t^i + P_t S_t^i = W_t + (D_t + P_t)S_{t-1}^i - \tau(P_t - P_{t-1})S_{t-1}^i + T_t^i \tag{76}$$

which is differentiable everywhere. Then, the subjective valuation becomes

$$Q_t^i = \delta \mathbb{E}^{\mathcal{P}^i}\left[D_{t+1} + P_{t+1} - \tau(P_{t+1} - P_t)\middle|\, z_t^i\right]$$

Equilibrium. An equilibrium with full specialization naturally emerges in the assumed configuration of beliefs without the need for requiring large amounts of disagreement, as there is not a no-trade region. Given $\pi^A_{H\to H} > \pi^B_{H\to H}$ and $\pi^A_{L\to L} > \pi^B_{L\to L}$, the B-REE with accrual taxes implies this system of two equilibrium conditions:

$$\begin{aligned} P_H &= \delta[\pi_A D_H + (1-\pi_A)D_L + (\pi_A P_H + (1-\pi_A)P_L)(1-\tau) + \tau P_H] \\ P_L &= \delta[\pi_B D_L + (1-\pi_B)D_H + (\pi_B P_L + (1-\pi_B)P_H)(1-\tau) + \tau P_L] \end{aligned}$$

Under the knife-edge condition $\pi_A = 1 - \pi_B$, expected dividends are identical in both states, $D^e = \pi_A D_H + (1-\pi_A)D_L = (1-\pi_B)D_H + \pi_B D_L$, so solving the two equations yields $P_H = P_L \equiv P^*$ with

$$P^* = \frac{\delta}{1-\delta} D^e$$

so there are no capital gains in equilibrium and the accrual tax rate τ is irrelevant in this case.

Consider the small deviation from this equilibrium as in the previous section. Using the subjective price model (22), the beliefs updating equation simplifies to

$$\beta_t = \beta_{t-1} + g\Big(P_{t-1} - \beta_{t-1}\Big) + g\varepsilon_t^b \tag{77}$$

where $g = \frac{\sigma_s^2}{\sigma_s^2 + \sigma_P^2}$ with $\sigma_s^2 = \frac{\sigma_\nu^2 + \sqrt{\sigma_\nu^4 + 4\sigma_\nu^2\sigma_P^2}}{2}$. It follows that the equilibrium price under Learning boils down to

$$P_t = \frac{\delta}{1-\delta\tau}(D^e + (1-\tau)\beta_t) \tag{78}$$

When agents learn about prices, prices remain state-independent but they inherit the beliefs dynamics.

Taxes and volatility. Taking variances on both sides of the equilibrium pricing function (78),

$$\text{Var}(P_t) = \omega \text{Var}(\beta_t) \tag{79}$$

where $\omega \equiv \left(\frac{\delta(1-\tau)}{1-\delta\tau}\right)^2$. Intuitively, the price variance depends on the variance of beliefs up to a scaling factor that determines the pass-through from beliefs to price volatility. Using this variance decomposition, the total derivative with respect to τ is given by:

$$\frac{d\text{Var}(P_t)}{d\tau} = \frac{d\omega}{d\tau}\text{Var}(\beta_t) + \omega \frac{d}{d\tau}\text{Var}(\beta_t) \tag{80}$$

As in the realization-based case, τ influences price volatility through two channels: the loading ω and the learning process that determines beliefs' dynamics. The following proposition characterizes the directional effect of the tax.

Proposition 4. *In the equilibrium with learning about prices with a capital gains tax on accrual, the loading, the variance of price beliefs, and the variance of the stock price are all decreasing in τ:*

$$\frac{d\omega}{d\tau} < 0, \quad \frac{d}{d\tau} Var(\beta_t) < 0, \quad \frac{d}{d\tau} Var\Big(P_t\Big) < 0$$

Proof. Appendix D.6.

The intuition of the proposition is straightforward: accrual-based taxation retains the pass-through effect entirely avoiding the lock-in effect. From a volatility point of view, that makes it appealing.

5.1.2. Financial Transactions Tax

Consider now a tax τ^F on the value of financial transactions. The budget constraint becomes

$$C_t^i + P_t S_t^i = W_t + (D_t + P_t) S_{t-1}^i - \tau^F |S_{t-1}^i - S_t^i| P_t + T_t^i \tag{81}$$

Proceeding exactly as in the CGT case, one obtains the following optimal demand function:

$$S_t^i = \begin{cases} \bar{S} & \text{if } Q_t^i \geq P_t(1+\tau^F) \\ 0 & \text{if } Q_t^i \leq P_t(1-\tau^F) \\ S_{t-1}^i & \text{if } P_t(1-\tau^F) \leq Q_t^i \leq P_t(1+\tau^F) \end{cases} \tag{82}$$

where

$$Q_t^i = \delta \mathbb{E}_t^{\mathcal{P}^i}\left[D_{t+1} + P_{t+1}(1 - \tau^F \mathbb{1}\{S_{t+1}^i < S_t^i\} + \tau^F \mathbb{1}\{S_{t+1}^i > S_t^i\})\right] \tag{83}$$

This delivers a symmetric no-trade band $[P_t(1-\tau^F), P_t(1+\tau^F)]$ of width $2\tau^F P_t$ around Q_t^i, in contrast with the CGT case where the sell threshold depends on the embedded gain via

the basis. Note that under an FTT the proportional trading wedge τ^F is invariant across states (it depends only on the transaction value), whereas under a CGT the effective selling wedge depends on the investor's basis through the embedded gain.

Equilibrium. For brevity, I first outline the B-REE in this context, and then expose the case when agents learn about prices. Before stating the equilibrium characterization, let me introduce compact notation. Temporarily, I write again the transition probabilities as $\pi^i_{s\to s}$ for $i \in \{A, B\}$ and $s \in \{H, L\}$. Define the scalar

$$M^F \equiv \Big((1+\tau^F) - \delta\pi^A_{H\to H}\Big)\Big((1+\tau^F) - \delta\pi^B_{L\to L}\Big) - \delta^2(1-\pi^A_{H\to H})(1-\pi^B_{L\to L})(1-\tau^F)^2$$

which is positive under the maintained parameter restrictions $\delta \in (0,1)$, $\pi^A_{H\to H}, \pi^B_{L\to L} \in [0,1]$, and $\tau^F \in [0,1]$.[44] Let

$$K^F \equiv \frac{\delta}{M^F}\begin{pmatrix} (1+\tau^F) - \delta\pi^B_{L\to L} & \delta(1-\pi^A_{H\to H})(1-\tau^F) \\ \delta(1-\pi^B_{L\to L})(1-\tau^F) & (1+\tau^F) - \delta\pi^A_{H\to H}\end{pmatrix}$$

Let $\mathbf{P} \equiv (P_H, P_L)' \in \mathbb{R}^2_+$ be an arbitrary candidate stationary state-price vector. Define the following price-dependent belief thresholds:

$$\begin{aligned}
\underline{\pi}^A_{H\to H} &\equiv \frac{(1+\delta)\,(1-\pi^B_{L\to L})\,(D_H - D_L) \;+\; \delta\,\pi^B_{L\to L}\,D_L}{\delta\,D_L \;+\; \big(1+2\delta(1-\pi^B_{L\to L})\big)\,(D_H - D_L)} \\
\bar{\pi}^B_{H\to H}(\mathbf{P}) &\equiv \frac{\frac{(1-\tau^F)P_H}{\delta} - \Big(D_L + (1+\tau^F)P_L\Big)}{(D_H - D_L) + (1-\tau^F)P_H - (1+\tau^F)P_L} \\
\bar{\pi}^A_{L\to H}(\mathbf{P}) &\equiv \frac{\frac{(1-\tau^F)(1-\delta)}{\delta}P_L - D_L}{(D_H - D_L) + (1+\tau^F)P_H - (1-\tau^F)P_L}
\end{aligned} \tag{84}$$

Notice the thresholds $\bar{\pi}^B_{H\to H}(\mathbf{P})$ and $\bar{\pi}^A_{L\to H}(\mathbf{P})$ are not primitive objects: they are functions of prices. In equilibrium, they are evaluated at the endogenous price vector $\mathbf{P}^*$ characterized in the proposition below. Similarly, define the following price-dependent feasibility bounds:

$$\bar{D}_L(\mathbf{P}) \;\equiv\; \min\Big\{\frac{(1-\tau^F)(1-\delta)}{\delta}\,P_L,\; \Big(\frac{1-\tau^F}{\delta}\Big)P_H - (1+\tau^F)P_L\Big\}$$

$$\bar{\tau}^F(\mathbf{P}) \;\equiv\; \frac{(D_H - D_L) + (P_H - P_L)}{P_H + P_L}$$

These bounds summarize sufficient conditions ensuring that the thresholds in (84) are well-defined (in particular, that the denominators are positive) and feasible in the probability

[44] Let $u = 1+\tau^F$ and $v = 1-\tau^F$. Then $M^F(a,b)$ is bilinear in $a = \pi^A_{H\to H}$ and $b = \pi^B_{L\to L}$, so its minimum on $[0,1]^2$ is attained at a corner; evaluating $(a,b) \in \{0,1\}^2$ yields $M^F \in \{u^2 - \delta^2 v^2,\; u(u-\delta),\; (u-\delta)^2\}$, all strictly positive for $\delta \in (0,1)$ and $\tau^F \in [0,1]$.

simplex when evaluated at $\mathbf{P} = \mathbf{P}^*$. With this notation in place, the following proposition characterizes the B–REE with a financial transaction tax.

Proposition 5. *Assume that beliefs satisfy:*

$$(i) \quad \pi^A_{H\to H} > \underline{\pi}^A_{H\to H} \tag{85}$$

$$(ii) \quad \pi^B_{H\to H} \le \bar{\pi}^B_{H\to H}(\mathbf{P}^*) \tag{86}$$

$$(iii) \quad \pi^A_{L\to H} \le \bar{\pi}^A_{L\to H}(\mathbf{P}^*) \tag{87}$$

and that the conditions $D_L \le \bar{D}_L(\mathbf{P}^*)$ *and* $\tau^F \le \bar{\tau}^F(\mathbf{P}^*)$ *hold. Then there exists a competitive equilibrium with*

$$(S^A_t, S^B_t) = \begin{cases} (1,0), & \text{if } s_t = H \\ (0,1), & \text{if } s_t = L \end{cases} \tag{88}$$

$$P_t = \begin{cases} P^*_H, & \text{if } s_t = H \\ P^*_L, & \text{if } s_t = L \end{cases} \tag{89}$$

where the equilibrium state-price vector $\mathbf{P}^*$ *is given by*

$$\mathbf{P}^* \;=\; K^F \mathbf{D}^e \tag{90}$$

Proof. Appendix D.7.

The proposition delivers the same qualitative price structure as in the CGT case: equilibrium prices are a linear function of expected dividends, with the projection matrix now incorporating the buy–sell wedge created by τ^F. The disagreement requirements are expressed through price-dependent belief thresholds. In state H, type B's cutoff governs taxed liquidation on an $L \to H$ arrival (entering with $S^B_{t-1} = 1$): she needs to be sufficiently pessimistic about staying in H, so that selling is optimal. Symmetrically, in state L, type A's cutoff governs taxed liquidation on an $H \to L$ arrival (entering with $S^A_{t-1} = 1$): she needs to be sufficiently pessimistic about returning to H, so selling is optimal. Condition (i) ensures $P^*_H > P^*_L$ under $D_H > D_L$. The feasibility bounds $D_L \le \bar{D}_L(\mathbf{P}^*)$ and $\tau^F \le \bar{\tau}^F(\mathbf{P}^*)$ are implicit sufficient restrictions, evaluated at $\mathbf{P}^*$, that guarantee these price-dependent thresholds are well-defined and admit solutions in $[0,1]$.[45]

With a financial transaction tax, trading on switches is sustainable only if beliefs are

[45] The feasibility restrictions on D_L and τ^F are technical sufficient conditions and are quite relaxed in typical parameterizations. For instance, for $\delta = 0.99$ and $\tau^F = 0.02$,

$$\frac{(1-\tau^F)(1-\delta)}{\delta} = \frac{0.98 \cdot 0.01}{0.99} \approx 9.9 \times 10^{-3},$$

so the first bound simply requires $D_L/P_L = O(1-\delta)$ and is slack whenever P_L reflects any meaningful

sufficiently dispersed, because buyers pay $(1+\tau^F)P_s^*$ while sellers receive only $(1-\tau^F)P_s^*$. To make this "required disagreement" explicit (as in the CGT case), rewrite the participation bounds (ii)–(iii) as wedge conditions. Define the high-state and low-state belief wedges at the equilibrium price vector $\mathbf{P}^*$ by

$$\kappa_H^F(\tau^F) \equiv \pi^A_{H\to H} - \bar{\pi}^B_{H\to H}(\mathbf{P}^*)$$

$$\kappa_L^F(\tau^F) \equiv \pi^A_{L\to H} - \bar{\pi}^A_{L\to H}(\mathbf{P}^*)$$

so that (ii) is equivalent to $\pi^B_{H\to H} \le \pi^A_{H\to H} - \kappa_H^F(\tau^F)$ and (iii) is equivalent to $\pi^A_{L\to H} \le \pi^A_{L\to H} - \kappa_L^F(\tau^F)$ (i.e., the relevant belief gaps must exceed κ_H^F and κ_L^F). The dependence on τ^F is explicit because $\bar{\pi}^B_{H\to H}(\mathbf{P})$ and $\bar{\pi}^A_{L\to H}(\mathbf{P})$ in (84) contain the net-of-tax trading prices $(1-\tau^F)P_H$, $(1+\tau^F)P_H$, $(1-\tau^F)P_L$, and $(1+\tau^F)P_L$; evaluated at $\mathbf{P}^*(\tau^F)$ from (90), this yields wedges $\kappa_H^F(\tau^F)$ and $\kappa_L^F(\tau^F)$ that are functions of τ^F both directly (through the $(1\pm\tau^F)$ terms) and indirectly (through the endogenous prices $\mathbf{P}^*(\tau^F)$).

Let again $\pi_A \equiv \pi^A_{H\to H}$ and $\pi_B \equiv \pi^B_{L\to L}$. Consider now the same small deviation from this equilibrium when agents have imperfect market knowledge. In this case, equilibrium prices satisfy this system of equations:

$$P_{H,t} = \frac{\delta}{1+\tau^F}\Big[\pi_A\big(D_H + \mathbb{E}_t^{\mathcal{P}}(P_{H,t+1})\big) + \big(1-\pi_A\big)\big(D_L + (1-\tau^F)\mathbb{E}_t^{\mathcal{P}}P_{L,t+1}\big)\Big] \tag{91}$$

$$P_{L,t} = \frac{\delta}{1+\tau^F}\Big[\pi_B\big(D_L + \mathbb{E}_t^{\mathcal{P}}(P_{L,t+1})\big) + \big(1-\pi_B\big)\Big(D_H + (1-\tau^F)\mathbb{E}_t^{\mathcal{P}}(P_{H,t+1})\Big)\Big] \tag{92}$$

The system comes from the relevant first order conditions and market clearing ($(1+\tau^F)P_H = Q_H^A$ and $(1+\tau^F)P_L = Q_L^B$). Using the same model of prices the was outlined for the CGT case, equilibrium prices can be compactly written as

$$\boldsymbol{P}_t^* = \frac{\delta}{1+\tau^F}(\boldsymbol{D}^e + \mathcal{B}^F\boldsymbol{\beta}_t) \tag{93}$$

where

$$\mathcal{B}^F = \begin{pmatrix} \pi_A & (1-\tau^F)\big(1-\pi_A\big) \\ (1-\tau^F)(1-\pi_B) & \pi_B \end{pmatrix} \tag{94}$$

The pricing equation (93) and the updating equation (25) fully characterize the dynamics of the model when agents learn about prices. As in the CGT case, equilibrium state prices become time-varying, driven by fluctuations in expected future prices. The

transition/resale option value. The second bound is slack if

$$\frac{P_H^*}{P_L^*} > \delta\frac{1+\tau^F}{1-\tau^F} = 0.99\cdot\frac{1.02}{0.98} \approx 1.03,$$

i.e., a spread above 3%. Ain the empirically relevant low-wedge region (e.g., $\tau^F \approx 0.02$) the restriction is comfortably satisfied.

FTT enters the pricing function through two distinct channels. First, it appears in the denominator, reflecting the transaction cost incurred upon purchase. Second, it enters the loading matrix $\mathcal{B}^F$, capturing the anticipated tax liability upon future sales: sellers internalize that realizing gains will trigger a transaction cost, which depresses the expected after-tax payoff.

Taxes and volatility. I start by analyzing the effects of the FTT on price volatility in the B–REE context. Let

$$\underline{\pi}_A^F \equiv 1 - \frac{1}{\delta}\left[\frac{1}{\delta(1-\pi_B)} - \frac{C_1\, C_0}{1-\tau^F C_1}\right]$$

with

$$C_0 \equiv (1-\delta)\Big(\frac{1}{\delta(1-\pi_B)} - 1\Big)$$

$$C_1 \equiv \frac{2\big((1+\tau^F) - \delta\pi_B\big) + 2\delta^2(1-\pi_B)^2(1-\tau^F)}{\big((1+\tau^F) - \delta\big)^2 - \delta^2(1-\pi_B)^2(1-\tau^F)^2}$$

Given this notation, the following proposition states the influence of a financial transaction tax on price volatility.

Proposition 6. *Consider the Bayesian–Rational Expectations Equilibrium with a Financial Transaction Tax characterized in Proposition 5. Additionally, suppose*

$$\textit{(i)} \quad \pi_A > \underline{\pi}_A^F \tag{95}$$

$$\textit{(ii)} \quad 1 - \tau^F C_1 > 0 \tag{96}$$

Then the variance of the stock price is increasing in τ^F:

$$\frac{d}{d\tau^F}\, Var(P_t^*) > 0$$

Proof. Appendix D.8.

As in the CGT case, a tax on financial transactions affects volatility through its effect on state-price dispersion. In the B–REE with an FTT, τ^F enters equilibrium prices through two conceptually distinct mechanisms. First, the purchase wedge $(1+\tau^F)^{-1}$ scales the intertemporal pricing kernel that capitalizes future payoffs. This force acts symmetrically across states and therefore tends to move P_H^* and P_L^* in the same direction. Second, the tax reduces the resale component of continuation values: whenever trade occurs, the resale term in the Euler equations is multiplied by $(1-\tau^F)$. This trading-option channel is inherently state-dependent because the resale value is weighted by the perceived probability of switching to the other state. Hence the transaction tax has a state-dependent effect whenever the equilibrium continuation value places different weight on resale across H

and L. The proof of Proposition 6 shows that the state-dependent trading-option channel dominates the common purchase-wedge scaling, so that an increase in τ^F depresses the low-state valuation relatively more than the high-state valuation and therefore increases dispersion $\Delta^F(\tau^F) \equiv P_H^* - P_L^*$.

Since, in this two-state model, price volatility is increasing in this wedge, the key to Proposition 6 is to ensure that dispersion increases with τ^F. The assumptions isolate exactly these requirements, conditional on existence of the B–REE set in Proposition 5. Assumption (i) is a convenient one-sided sufficient condition ensuring that the state-dependent trading-option channel dominates the common purchase-wedge scaling, so that the net positive sign emerges. Assumption (ii) is a purely technical sign restriction used in the proof: it ensures that the rearrangement delivering the one-sided bound $\underline{\pi}_A^F$ preserves inequality direction (equivalently, it rules out regions where the sufficient condition in terms of π_A would flip or become vacuous). Under these conditions, dispersion is positive and increasing in τ^F, implying $\frac{d}{d\tau^F}\mathrm{Var}(P_t^*) > 0$.

In the CGT economy, the realization wedge selects frontier beliefs in a way that tilts valuations toward the high state, raising P_H^* relative to P_L^* and thereby reinforcing state-price dispersion. In contrast, under an FTT the belief adjustments needed to sustain trade need not tilt dispersion in a single direction. When a trading constraint binds, a higher τ^F lowers the seller's net-of-tax proceeds in the relevant state; to keep the marginal seller willing to trade, equilibrium beliefs on the binding frontier must assign sufficiently high probability to being in the high state next period, so that the continuation value is large enough to compensate the seller for tax payments. Because this compensation can operate through either frontier, it can raise P_H^* (when trade is disciplined by the high-state constraint) or raise P_L^* (when it is disciplined by the low-state constraint). As a result, the frontier contribution to the dispersion wedge $\Delta^F(\tau^F) \equiv P_H^* - P_L^*$ is in general ambiguous.

Consider now the equilibrium with learning about prices. Beliefs dynamics are characterized by an equation like (33) but adjusting the loadings matrix. Naturally, the same logic of lemma 3 carries over to the case with FTT.

Lemma 4. *Consider the equilibrium with learning about prices with a Financial Transactions Tax, with pricing function (93) and belief updating (25). Suppose symmetric Kalman gains $g_H = g_L \equiv g > 0$ and adopt a mean-coefficient approximation to belief dynamics. Let $\Lambda^F \equiv \frac{\delta}{1+\tau^F}\mathcal{B}^F$. Then there exists $\bar{g} > 0$ such that, for all $g \in (0, \bar{g})$ and for each state $s \in \{H, L\}$:*

1. *The tax reduces the effective loadings,*

$$\frac{d\Lambda_s^F}{d\tau^F}\,\Sigma_\beta\,(\Lambda_s^F)' \;\leq\; 0$$

2. *The tax reduces belief fluctuations,*

$$\Lambda_s^F \frac{d\Sigma_\beta}{d\tau^F} (\Lambda_s^F)' \ < \ 0$$

3. *Up to a first-order approximation around the mean price, the tax reduces the variance of prices*

$$\frac{d}{d\tau^F} Var(P_{st}) \ < \ 0$$

Proof. Appendix D.9.

The lemma is the FTT analogue of Lemma 3 for the CGT economy, but the sign logic is simpler since τ^F reduces all elements of the loadings matrix. Hence, the pass-through from beliefs to prices is unambiguously attenuated, so prices are mechanically less sensitive to belief shocks. By the same token, smaller price forecast errors imply smaller belief revisions and therefore lower belief volatility under learning. In contrast, under a CGT, the basis-price channel generates an offsetting force: higher τ dampens some components of the belief-to-price map but raises others through the state-contingent realization wedge. As a result, the CGT counterpart requires additional restrictions to sign the net effect. Since the FTT effect is one-directional, the lemma imposes no further parameter restrictions beyond the maintained small-gain/mean-coefficient approximation, and is therefore more general than its CGT analogue.

Combining Proposition 6 and Lemma 4 one gets a theoretically ambiguous net effect, paralleling the CGT analysis. Using the law of total variance as in Proposition 3, the total derivative of the variance of log prices with respect to the FTT reads as

$$\frac{d\text{Var}(P_t^*)}{d\tau^F} = \mu(1-\mu) \underbrace{\frac{d}{d\tau^F}\left[P_H^* - P_L^*\right]^2}_{>0} + \underbrace{\frac{d}{d\tau^F}\Big[\mu \ \text{Var}(P_{Ht}) + (1-\mu)\text{Var}(P_{Lt})\Big]}_{<0} \tag{97}$$

The sign is theoretically ambiguous, with different components pulling in opposite directions; determining which dominates is an empirical question. This ambiguity may help rationalize the mixed findings in the empirical literature, as the two channels may operate with different intensities across institutional settings and market conditions. The framework also provides a lens through which to interpret the quantitative results of Adam et al. (2015), who find that an FTT stabilizes markets by reducing the magnitude and duration of stock price booms, yet destabilizes by increasing the frequency of boom-bust episodes. Within my framework, the former reflects the pass-through effect while the latter reflects a lock-in effect. Their conclusion that FTTs are mildly destabilizing overall is thus contingent on the relative strength of these opposing forces under their particular parameterization.

5.2. Quantitative Experiments

5.2.1. Experiment 1: Replacing Realization-Based by Accrual-Based Taxation

This experiment implements a regime change in capital gains taxation from a realization to an accrual base. As argued above, accrual-based taxation is appealing for a simple reason: it retains the stabilizing pass-through effect while eliminating the destabilizing lock-in effect.

Consider the asymmetric accrual regime with carryover losses described in Section 4.1.1. The experiment constructs a sequence of accrual-based CGT rates $\{\tau_t^A\}_{t=1}^T$ that would have generated the same revenues as the model-implied realization-based fiscal revenues in every period. Specifically, I follow these steps:

1. Use the baseline estimated model to compute a series of tax revenues under realization-based taxation. Total tax revenue is given by:

$$\mathcal{T}_t = \tau_t^{CG} \sum_i \mu^i \max\Big\{ \max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i) - L_t^i S_{t-1}^i, 0 \Big\} \tag{98}$$

2. Using the baseline parameters, simulate the model under accrual taxation. Revenue under accrual is given by:

$$\mathcal{T}_t^A = \tau_t^A \sum_i \mu^i \max\{0, P_t - P_{t-1} - L_t^i\} S_{t-1}^i \tag{99}$$

3. Find τ_t^A such that $\mathcal{T}_t^A = \mathcal{T}_t$. This requires solving the following equation each period:[46]

$$\tau_t^A = \frac{\mathcal{T}_t}{\sum_i \mu^i \max\{0, P_t(\tau_t^A) - P_{t-1} - L_t^i\} S_{t-1}^i} \tag{100}$$

4. Simulate the baseline model replacing $\{\tau_t^{CG}\}_{t=1}^T$ with $\{\tau_t^A\}_{t=1}^T$.

Figure 5 plots the time series for τ_t^A. The average across Monte Carlo simulations and time periods is 7.18%, which is substantially below the average effective rate in the data, 10.77%.[47] The lower rate reflects the broader tax base under accrual: because all gains

[46] The contemporaneous accrual base

$$B_t(\tau) \equiv \sum_i \mu^i \max\{0,\ P_t(\tau) - P_{t-1} - L_t^i\}\, S_{t-1}^i$$

can be zero in downturn-like periods. In that case, $\mathcal{T}_t^A(\tau) \equiv \tau\, B_t(\tau) = 0$ for all $\tau \geq 0$, so exact within-period revenue matching is mechanically infeasible unless $\mathcal{T}_t$ is also essentially zero. In simulations, I therefore implement the following rule: if $B_t(\tau) \leq \varepsilon_B$ for all $\tau \in [\tau_{\min}, \tau_{\max}]$, I set $\tau_t^A = \tau_{\min}$ (so $\mathcal{T}_t^A = 0$) and record the residual $\mathcal{T}_t^A - \mathcal{T}_t$ for diagnostics; otherwise, I solve $\mathcal{T}_t^A(\tau_t^A) = \mathcal{T}_t$. I find that periods in which the asymmetric accrual base collapses coincide with negligible realization-based capital-gains tax revenues, so the residual is tiny (in particular during episodes like the Great Recession).

[47] For each Monte Carlo replication $m = 1, \ldots, M$ I draw a sequence of idiosyncratic wage shocks

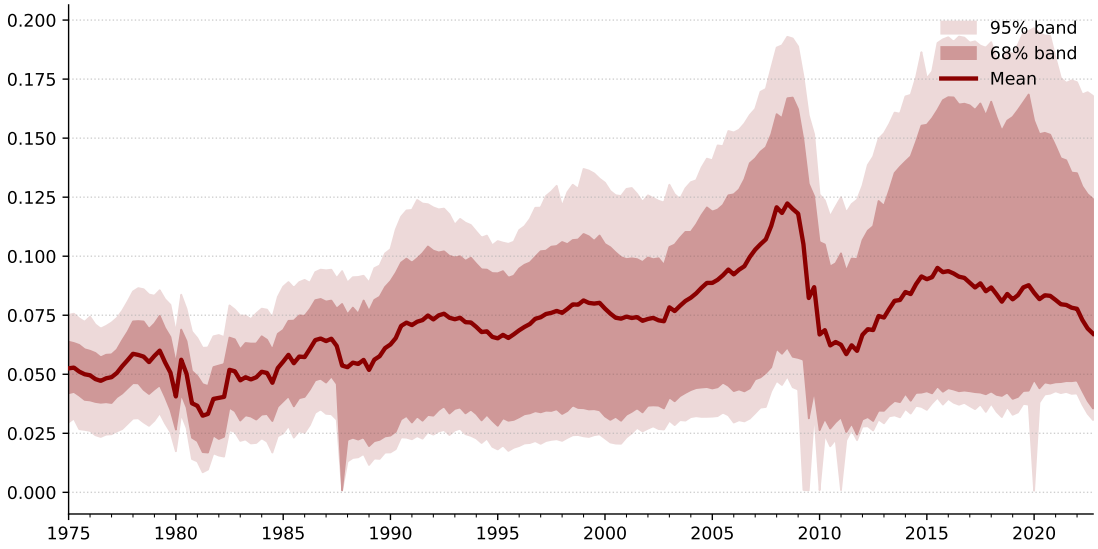


Figure 5: Revenue-neutral accrual-based capital gains tax. *The figure plots the time series of accrual-based CGT rates τ_t^A that would generate the same revenue as the estimated realization-based system. The solid line shows the mean across Monte Carlo simulations; shaded bands indicate the 68% and 95% percentile bands.*

are taxed regardless of realization decisions, a smaller rate suffices to generate equivalent revenue.[48]

Holding all else equal, the model implies that this regime change would have reduced valuations volatility by approximately 20% and returns volatility by about 10%. Additionally, the attenuation of shock propagation, as measured by the IRF coefficient, is substantial—roughly 40%—and the association between price changes and trading volume is cut by half. Thus, a lower tax applied in a less distortionary manner appears to pay off in terms of stock market stability.

5.2.2. Experiment 2: Financial Transaction Taxes versus Accrual-Based Capital Gains Taxes

In this experiment, I retain the realization-based CGT regime and introduce small supplementary taxes on either stock market transactions or unrealized gains with carryover losses. The technical details for incorporating these taxes into the quantitative framework

and compute the corresponding equilibrium outcomes and the revenue-neutral rate $\tau_t^{A,(m)}$ in each period. The solid line plots the cross-replication mean $\frac{1}{M}\sum_m \tau_t^{A,(m)}$. Shaded regions are percentile bands across replications (68%: 16–84; 95%: 2.5–97.5) and thus summarize the model-implied dispersion in τ_t^A driven by idiosyncratic income risk, holding the aggregate history fixed. These bands are not statistical confidence intervals and do not reflect parameter-estimation uncertainty.

[48]The realization-based rate satisfies $\mathcal{T}_t = \tau_t^{CG}\mathcal{G}_t$, so low-frequency movement in τ_t^{CG} largely reflects time variation in the realized-gains base $\mathcal{G}_t$ driven by asset-price growth and cycles. Under accrual taxation, τ_t^A solves $\mathcal{T}_t = \tau_t^A \widetilde{\mathcal{G}}_t(\tau_t^A)$, where the accrual base $\widetilde{\mathcal{G}}_t(\cdot)$ is endogenous to equilibrium prices. This fixed-point breaks any mechanical mapping from a declining τ_t^{CG} to a declining τ_t^A

of Section 3 are provided in Appendix E.

The motivation is twofold. First, financial transaction taxes have been routinely invoked as a potential tool to curb excess volatility (e.g., L. H. Summers and V. P. Summers (1989); Stiglitz (1989)) on the grounds that they would discourage speculative trading. The modern theoretical evidence, however, is mixed. While Dávila (2023) derive an ambiguous effect, Buss and Dumas (2019) show that a transaction tax can raise return volatility. Quantitative studies likewise suggest that FTTs need not reduce—and may even increase—volatility (e.g., Adam et al. (2015); Buss et al. (2016)). Here, I revisit this question leveraging the strong quantitative performance of the model. Second, an additional tax on unrealized gains—as distinct from a full regime change from realization to accrual—featured prominently in recent U.S. policy proposals under the Biden-Harris administration. While curbing excess volatility was not the primary stated objective, the analysis here offers an informative complementary perspective on the potential benefits of such a measure.

For the experiment, I simulate the model with small tax rates. FTT proposals in L. H. Summers and V. P. Summers (1989) have suggested rates around 0.5%, and the baseline optimal tax in Dávila (2023) is 0.37%. In their experiments, Adam et al. (2015) also consider higher rates. To capture the different magnitudes, I use $\tau^F \in \{0.5\%, 1\%, 2\%, 5\%\}$. The same set of rates is applied for the accrual-based CGT τ^A.

Effects of a small FTT. For small transaction taxes ($\tau^F \leq 0.02$), the FTT has a visible effect on the volatility of valuation ratios in levels but a muted effect on return volatility. As shown in Table 7, $\mathbb{V}(\ln P_t/D_t)$ falls from 19.22 at $\tau^F = 0.005$ to 12.88 at $\tau^F = 0.02$. To understand the underlying drivers, I use the following variance identity, which follows from the variance of a first difference:[49]

$$\text{Var}(\ln PD_t) = \frac{\text{Var}(\Delta \ln(PD_t))}{2(1-\rho)} \tag{101}$$

where $\rho \equiv \text{corr}(\ln PD_t, \ln PD_{t-1})$. Equation (101) shows that $\text{Var}(\ln PD_t)$ can fall either because the variance of valuation changes $\text{Var}(\Delta \ln PD_t)$ declines (smaller innovations) or because persistence falls (larger $1-\rho$). In the table, over $\tau^F \in \{0.005, 0.01, 0.02\}$, $\text{Var}(\Delta \ln PD_t)$ is essentially unchanged, while the serial correlation declines from 98.25 to 97.21. Thus, the reduction in the volatility of valuation ratios is driven primarily by a decline in their persistence rather than by period-to-period fluctuations.

Consistent with this diagnostic, return volatility barely responds to small FTT rates, since return volatility is primarily driven by the variance of valuation changes rather than by the persistence of the valuation ratio. The annualized $\sigma(\ln R^s_t)$ remains around 15: it

[49] Let $x_t \equiv \ln(P_t/D_t)$. By the variance formula, $\mathbb{V}(x_t - x_{t-1}) = \mathbb{V}(x_t) + \mathbb{V}(x_{t-1}) - 2\text{Cov}(x_t, x_{t-1})$. Under covariance stationarity, $\mathbb{V}(x_t) = \mathbb{V}(x_{t-1}) \equiv \mathbb{V}(x)$, and $\text{Cov}(x_t, x_{t-1}) = \rho\,\mathbb{V}(x)$. Hence $\mathbb{V}(\Delta x_t) = 2\mathbb{V}(x)(1-\rho)$, which implies (101).

declines only from 15.12 at $\tau^F = 0.005$ to 14.77 at $\tau^F = 0.02$. Moreover, the FTT amplifies shock propagation: the one-year response of valuations to a price-growth shock rises from 2.69 to 3.01 over the same range. Hence, at low rates the FTT alters the dynamics of the price-dividend ratio in a way that lowers unconditional dispersion, but it does not consistently attenuate the short-run movements in valuations that are most closely tied to return volatility.

Effects of small accrual-based capital gains taxes. The accrual-based capital gains tax operates through a fundamentally different mechanism. As τ^A increases, it delivers a broad-based reduction in volatility that is visible both in valuation innovations and in returns. $\mathrm{Var}(\Delta \ln PD_t)$ falls sharply—from 0.51 at $\tau^A = 0.005$ to 0.32 at $\tau^A = 0.02$, and to 0.18 at $\tau^A = 0.05$—and annualized return volatility declines in tandem, from 13.85 to 10.88 and then to 8.39. At the same time, the serial correlation of $\ln P_t/D_t$ rises (from 98.58 to 99.08 to 99.44), indicating that valuations become both smoother at high frequencies and more persistent at low frequencies. The IRF statistic corroborates a genuine dampening of shock transmission: the one-year valuation response falls monotonically with τ^A (from 2.48 at 0.005 down to 2.10 at 0.02). Taken together, these patterns are consistent with CGT-A acting directly on the marginal incentive to generate short-run price movements, thereby reducing both the variance of valuation changes and the propagation of price-growth shocks into valuations.

Additional moments. Two additional moments reinforce the contrast in mechanisms and highlight important side effects. First, the comovement of trading volume with absolute price changes, $\mathrm{corr}(TV_t, |P_t/P_{t-1} - 1|)$, is essentially unchanged under small FTT rates (approximately 52 percent) but declines systematically under CGT-A (from 48.0 to 37.7 as τ^A rises from 0.005 to 0.02). This suggests that CGT-A weakens the link between trading intensity and large price movements, whereas the FTT does not. Second, the FTT substantially raises mean valuations—$\mathbb{E}(PD_t)$ climbs from 49.05 to 61.56 by $\tau^F = 0.02$—whereas CGT-A leaves mean valuations broadly stable or slightly lower (46.76 to 45.52). Thus, CGT-A achieves lower fluctuations without affecting mean valuations much.[50]

Behavior at higher rates. At the higher rate $\tau^F = 0.05$, the FTT becomes clearly destabilizing: $\mathrm{Var}(\ln PD_t)$ jumps to 41.46, $\mathrm{Var}(\Delta \ln PD_t)$ rises to 0.90, $\sigma(\ln R_t^s)$ increases to 18.70, and the IRF more than doubles to 6.99. At this level, FTT sharply reduces the association between absolute price movements and trading volume. In this regime, the mechanism is a lock-in margin: the tax discourages rebalancing so that turnover becomes

[50] Notice that while the analytic results in Section 5.1.2 the effect of τ^F on equilibrium prices is negative, the effect here goes in the opposite direction. This reflects the influence of other forces, such as risk-aversion, and general equilibrium effects. Likewise, Adam et al. (2015) find that transaction taxes raise mean valuations.

Table 7: Selected statistics under a Financial Transaction Tax and an Accrual-Based Capital Gains Tax. *This table reports model-implied moments for alternative tax rates. Moments are defined as in Table 4, except two additional moments $Var(\Delta lnPD_t)$ and $corr(lnPD_t, lnPD_{t-1})$ which are the variance of the log price-dividend ratio growth and the serial correlation of the log price-dividend ratio. The case without additional taxes, that is, $\tau^F = 0$ for the Financial Transaction Tax and to $\tau^A = 0$ for the Accrual-Based Capital Gains Tax, corresponds to the baseline results in Table 4.*

	FTT (τ^F)				CGT-A (τ^A)			
	0.5%	1%	2%	5%	0.5%	1%	2%	5%
$\mathbb{E}(PD_t)$	49.05	51.73	61.56	77.78	46.76	46.21	45.52	43.98
$\mathbb{E}(r_t^s)$	8.48	8.14	6.99	6.00	8.55	8.40	8.23	8.09
$\text{Var}(\ln PD_t)$	19.22	17.96	12.88	41.46	19.72	19.16	18.70	17.86
$\text{Cov}(\ln PD_t, \bar{r}_t)$	-14.32	-13.15	-8.94	-32.98	-14.58	-13.89	-13.15	-11.92
$\text{Cov}(\ln PD_t, \bar{d}_t)$	4.90	4.80	3.93	8.48	5.14	5.27	5.55	5.95
$\text{Var}(\Delta \ln PD_t)$	0.61	0.58	0.59	0.90	0.51	0.42	0.32	0.18
$\text{corr}(\ln PD_t, \ln PD_{t-1})$	98.25	98.16	97.21	98.76	98.58	98.80	99.08	99.44
$\sigma(\ln R_t^s)$	15.12	14.79	14.77	18.70	13.85	12.58	10.88	8.39
$\mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 1, I_t) - \mathbb{E}(PD_{t+4}\|\varepsilon_t^P = 0, I_t)$	2.69	2.73	3.01	6.99	2.48	2.30	2.10	1.74
$\text{corr}(TV_t, \lvert P_t/P_{t-1} - 1\rvert)$	51.90	52.16	52.15	32.77	48.00	44.01	37.66	24.46

muted and less responsive precisely when prices move the most, while price responses and valuation volatility are amplified. Over the same range, CGT-A continues to move all moments smoothly and monotonically in the stabilizing direction; there, the weaker comovement between trading and prices reflects smaller price fluctuations (pass-through stabilization) rather than suppressed trading.

In summary, CGT-A appears both more targeted at the sources of return volatility and more robust to the precise choice of the tax rate, while the FTT can generate adverse dynamics once it is set even moderately high. These findings suggest that policymakers seeking to reduce excess volatility through the tax system should favor instruments that operate on the incentive to hold appreciated assets rather than on the act of trading itself.

6. Conclusions

Returning to a central question in the theory of capital gains taxation, this paper develops a framework to analyze how capital gains taxes affect asset price volatility. The model incorporates key features of real-world tax systems that have proved challenging to incorporate in equilibrium models: capital gains are taxed only upon realization and asymmetrically—gains are taxed, while losses are carried forward to offset future gains. Within this environment, capital gains taxation generates two opposing forces. A lock-in channel widens the no-trade region and can amplify price jumps; a pass-through chan-

nel dampens the transmission of beliefs fluctuations into prices, which is at the core of the price-belief feedback loop that generates excess volatility. The net effect is therefore theoretically ambiguous and must be resolved empirically.

I apply the framework to evaluate the consequences of the sequence of U.S. capital gains tax cuts since the 1970s. An estimated version of the model matches a comprehensive set of moments remarkably well and implies that, over this historical episode, tax cuts increased volatility. Quantitatively, the model attributes to these tax cuts a 35% increase in price–dividend ratio volatility. Volatility rises by 50% because lower capital-gains taxation strengthens the pass-through from beliefs to prices, amplifying self-fulfilling fluctuations. The resulting attenuation of the lock-in channel—together with general-equilibrium feedbacks—is not strong enough to offset this amplification. Similarly, the model predicts a 30% increase in stock return volatility, aligned with the econometric evidence in Z. Dai, Shackelford, and Zhang (2013) and providing a structural interpretation of their reduced-form estimates.

Leveraging the model's quantitative performance, I use it to study alternative tax arrangements. The main policy lesson is that adding a small tax on unrealized gains can improve financial stability: it strengthens the stabilizing pass-through mechanism while avoiding the destabilizing lock-in channel. By contrast, a financial transactions tax delivers mixed results and, if miscalibrated, can meaningfully distort trading and price dynamics, potentially making markets more fragile rather than more stable.

The analysis in this paper is deliberately positive. A natural next step is to study optimal policy in this richer environment. While Sahm (2008) and Aguiar, Moll, and Scheuer (2024) analyze optimal policy when comparing realization- and accrual-based regimes, the former relies on a highly stylized setting and the latter abstracts from the lock-in effect. Revisiting this debate in a framework that nests both lock-in and pass-through forces could yield new insights into the theory of capital gains taxation.

The model is an exchange economy, so it abstracts from feedback from asset prices to the macroeconomy. Introducing production would allow to study how tax-induced volatility transmits to investment and growth. It would also enable an integrated analysis of personal and corporate capital taxation, providing a more complete view of how lower capital taxes—both personal and corporate—shape macroeconomic fluctuations.

The model focuses on the aggregate stock market to derive an aggregate effect of tax cuts consistent with the available cross-sectional evidence. A natural extension is to introduce multiple risky assets that differ in their exposure to capital gains taxation. If tax wedges are asset-specific, capital gains taxes can affect not only volatilities but also return correlations and hedging properties, with implications for equilibrium risk premia, portfolio allocation, and shock propagation across markets. I leave these cross-asset comovement effects for future research.

Although the paper already engages with the cross-sectional evidence in Z. Dai, Shack-

elford, and Zhang (2013), a natural complementary source of cross-sectional variation is international. Anecdotal accounts suggest that tax reductions have been followed by episodes of financial stress and boom–bust dynamics, including the Scandinavian crisis and housing booms in Spain, Denmark, and Ireland. Conversely, Japan's extraordinary run-up in equity prices during the 1980s occurred in a largely tax-free environment, and the market's peak coincided with the introduction of a tax on stock capital gains in April 1989. While systematically exploiting this cross-country and historical variation is challenging—because data quality, institutional heterogeneity, general-equilibrium spillovers, and cross-border capital flows complicate identification—it remains a promising research avenue for assessing the role of capital gains taxation in financial instability globally.

Methodologically, the paper proposes a tractable way to solve and estimate rich heterogeneous agent models under an information structure that relaxes full-information rational expectations in a manner consistent with survey evidence. I have illustrated the approach in a highly nonlinear setting; a natural extension is to embed it in environments with a richer financial sector, endogenous leverage, and production in order to study economic crises and to evaluate the joint design of macroprudential and monetary policy. One conjecture is that capital taxation used for financial-stability purposes can provide an additional policy margin, potentially easing potential conflicts between macroeconomic and financial stability when using a single instrument -the interest rate.[51]

Economists have long sought tax instruments that can mitigate financial instability. Although an FTT has been a prominent proposal, its implications for volatility are mixed. Against this backdrop, this paper shows that capital gains taxation can serve as an effective tool to curb excessive asset price volatility.

[51] See, for instance, this speech by MinnFed president. https://www.minneapolisfed.org/article/2017/monetary-policy-and-bubbles.

References

Adam, Klaus (2020). *Monetary policy challenges from falling natural interest rates*. Tech. rep. ECB conference.

Adam, Klaus, Johannes Beutel, Albert Marcet, and Sebastian Merkel (2015). "Can a financial transaction tax prevent stock price booms?" In: *Journal of monetary economics* 76, S90–S109.

Adam, Klaus and Albert Marcet (2011). "Internal rationality, imperfect market knowledge and asset prices". In: *Journal of Economic Theory* 146.3, pp. 1224–1252.

Adam, Klaus, Albert Marcet, and Johannes Beutel (2017). "Stock price booms and expected capital gains". In: *American Economic Review* 107.8, pp. 2352–2408.

Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini (2016). "Stock market volatility and learning". In: *The Journal of Finance* 71.1, pp. 33–82.

Adam, Klaus and Sebastian Merkel (2019). "Stock price cycles and business cycles". In.

Agersnap, Ole and Owen Zidar (2021). "The tax elasticity of capital gains and revenue-maximizing rates". In: *American Economic Review: Insights* 3.4, pp. 399–416.

Aguiar, Mark, Benjamin Moll, and Florian Scheuer (2024). "Putting the 'Finance'into 'Public Finance': A Theory of Capital Gains Taxation". In.

Anagnostopoulos, Alexis, Eva Cárceles-Poveda, and Danmo Lin (2012). "Dividend and capital gains taxation under incomplete markets". In: *Journal of Monetary Economics* 59.7, pp. 599–611.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2018). "Extrapolation and bubbles". In: *Journal of Financial Economics* 129.2, pp. 203–227.

Becchetti, Leonardo, Massimo Ferrari, and Ugo Trenta (2014). "The impact of the French Tobin tax". In: *Journal of Financial Stability* 15, pp. 127–148.

Belda, Pau, Janko Heineken, and Adrian Ifrim (2025). "Heterogeneous Beliefs and Wealth Inequality". In: *Working Paper*.

Benhabib, Jess, Wei Cui, and Jianjun Miao (2024). "Capital income jumps and wealth distribution". In: *Quantitative Economics* 15.4, pp. 1197–1247.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2024). "Belief overreaction and stock market puzzles". In: *Journal of Political Economy* 132.5, pp. 1450–1484.

Brun, Lídia and Ignacio González (2017). "Tobin's Q and Inequality". In: *Available at SSRN 3069980*.

Buss, Adrian and Bernard Dumas (2019). "The dynamic properties of financial-market equilibrium with trading fees". In: *The Journal of Finance* 74.2, pp. 795–844.

Buss, Adrian, Bernard Dumas, Raman Uppal, and Grigory Vilkov (2016). "The intended and unintended consequences of financial-market regulations: A general-equilibrium analysis". In: *Journal of Monetary Economics* 81, pp. 25–43.

Campbell, John Y and John H Cochrane (1999). "By force of habit: A consumption-based explanation of aggregate stock market behavior". In: *Journal of political Economy* 107.2, pp. 205–251.

Cappelletti, Giuseppe, Giovanni Guazzarotti, and Pietro Tommasino (2017). "The stock market effects of a securities transaction tax: quasi-experimental evidence from Italy". In: *Journal of Financial Stability* 31, pp. 81–92.

Cloe, Carl W (1952). "Capital gains and the changing price level". In: *National Tax Journal* 5.3, pp. 207–217.

Cochrane, John H (2009). *Asset pricing: Revised edition*. Princeton university press.

Constantinides, George M (1983). "Capital market equilibrium with personal tax". In: *Econometrica: Journal of the Econometric Society*, pp. 611–636.

Dai, Min, Hong Liu, Chen Yang, and Yifei Zhong (2015). "Optimal tax timing with asymmetric long-term/short-term capital gains tax". In: *The Review of Financial Studies* 28.9, pp. 2687–2721.

Dai, Zhonglan, Douglas Shackelford, and Harold Zhang (2013). "Capital gains taxes and stock return volatility". In: *The Journal of the American Taxation Association* 35.2, pp. 1–31.

Dávila, Eduardo (2023). "Optimal financial transaction taxes". In: *The Journal of Finance* 78.1, pp. 5–61.

Deng, Yongheng, Xin Liu, and Shang-Jin Wei (2018). "One fundamental and two taxes: When does a Tobin tax reduce financial price volatility?" In: *Journal of Financial Economics* 130.3, pp. 663–692.

Dowd, Tim, Robert McClelland, and Athiphat Muthitacharoen (2015). "New evidence on the tax elasticity of capital gains". In: *National Tax Journal* 68.3, pp. 511–544.

Ehling, Paul, Michael Gallmeyer, Sanjay Srivastava, Stathis Tompaidis, and Chunyu Yang (2010). "Portfolio choice with capital gain taxation and the limited use of losses". In: *Available at SSRN 1003034*.

Ehling, Paul, Stathis Tompaidis, and Chunyu Yang (2019). "Tax Collection from Realized Capital Gains on Equity". In: *Available at SSRN 3349058*.

Eichfelder, Sebastian, Mona Noack, and Felix Noth (2022). "The impact of financial transaction taxes on stock markets: Short-run effects, long-run effects, and reallocation of trading activity". In: *National Tax Journal* 75.3, pp. 539–569.

Feenberg, Daniel and Elisabeth Coutts (1993). "An introduction to the TAXSIM model". In: *Journal of Policy Analysis and management* 12.1, pp. 189–194.

Fernández-Villaverde, Jesús (2025). *Deep Learning for Solving Economic Models*. Tech. rep. National Bureau of Economic Research.

Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuno (2023). "Financial frictions and the wealth distribution". In: *Econometrica* 91.3, pp. 869–901.

Gemmill, Robert F (1956). "The Effect of the Capital Gains Tax on Asset Prices". In: *National Tax Journal* 9.4, pp. 289–301.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus (2021). "Five facts about beliefs and portfolios". In: *American Economic Review* 111.5, pp. 1481–1522.

Gravelle, Jane G (2020). "Capital gains tax options: Behavioral responses and revenues". In: *CRS Report for Congress*. Vol. 41364.

Hakansson, Nils H. (1970). "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions". In: *Econometrica* 38.5, pp. 587–607.

Han, Jiequn, Yucheng Yang, et al. (2021). "Deepham: A global solution method for heterogeneous agent models with aggregate shocks". In: *arXiv preprint arXiv:2112.14377*.

Haugen, Robert A and A James Heins (1969). "The effects of the personal income tax on the stability of equity value". In: *National Tax Journal* 22.4, pp. 466–471.

Haugen, Robert A and Dean W Wichern (1973). "The diametric effects of the capital gains tax on the stability of stock prices". In: *The Journal of Finance* 28.4, pp. 987–996.

Kase, Hanno, Leonardo Melosi, and Matthias Rottner (2022). *Estimating nonlinear heterogeneous agents models with neural networks*. Centre for Economic Policy Research.

Klein, Peter (1999). "The capital gain lock-in effect and equilibrium returns". In: *Journal of Public Economics* 71.3, pp. 355–378.

Kueng, Lorenz (2014). *Tax news: The response of household spending to changes in expected taxes*. Tech. rep. National Bureau of Economic Research.

Lettau, Martin, Sydney C Ludvigson, and Jessica A Wachter (2008). "The declining equity premium: What role does macroeconomic risk play?" In: *The Review of Financial Studies* 21.4, pp. 1653–1687.

Marcet, Albert (1988). "Solving nonlinear stochastic growth models by parametrizing expectations." In.

Marcet, Albert and Kenneth J Singleton (1999). "Equilibrium asset prices and savings of heterogeneous agents in the presence of incomplete markets and portfolio constraints". In: *Macroeconomic Dynamics* 3.2, pp. 243–277.

McGrattan, Ellen R (2023). "Taxes, regulations, and the value of US corporations: A reassessment". In: *Review of Economic Dynamics* 50, pp. 131–145.

McGrattan, Ellen R and Edward C Prescott (2003). "Average debt and equity returns: Puzzling?" In: *American Economic Review* 93.2, pp. 392–397.

— (2005). "Taxes, Regulations, and the Value of US and UK Corporations". In: *The Review of Economic Studies* 72.3, pp. 767–796.

Milgrom, Paul and Ilya Segal (2002). "Envelope theorems for arbitrary choice sets". In: *Econometrica* 70.2, pp. 583–601.

Moll, Benjamin (2025). "The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics". In: *The Economic Journal*, ueaf104.

Rosenthal, Steve and Lydia Austin (2016). "The dwindling taxable share of US corporate stock". In: *Tax Notes* 151.6.

Sahm, Marco (2008). "Methods of capital gains taxation and the impact on asset prices and welfare". In: *National Tax Journal* 61.4, pp. 743–768.

Sarin, Natasha, Lawrence Summers, Owen Zidar, and Eric Zwick (2022). "Rethinking how we score capital gains tax reform". In: *Tax Policy and the Economy* 36.1, pp. 1–33.

Sialm, Clemens (2009). "Tax changes and asset pricing". In: *American Economic Review* 99.4, pp. 1356–83.

Somers, Harold M (1948). "An economic analysis of the capital gains tax". In: *National Tax Journal* 1.3, pp. 226–232.

— (1960). "Reconsideration of the capital gains tax". In: *National Tax Journal* 13.4, pp. 289–309.

Stiglitz, Joseph E (1983). "Some aspects of the taxation of capital gains". In: *Journal of Public Economics* 21.2, pp. 257–294.

— (1989). "Using tax policy to curb speculative short-term trading". In: *Regulatory Reform of Stock and Futures Markets: A Special Issue of the Journal of Financial Services Research*, pp. 3–17.

Summers, Lawrence H and Victoria P Summers (1989). "When financial markets work too well: A cautious case for a securities transactions tax". In: *Journal of financial services research* 3.2, pp. 261–286.

Viard, Alan D (2000). "Dynamic asset pricing effects and incidence of realization-based capital gains taxes". In: *Journal of Monetary Economics* 46.2, pp. 465–488.

Xiong, Wei (2013). "Bubbles, crises, and heterogeneous beliefs". In: *Handbook of Systemic Risk* Chapter 24, pp. 663–713.

Appendix

A. Data Sources

Stock market data. Stock prices, dividends and CPI inflation come from Shiller's website. The monthly data has been transformed into quarterly frequency by taking the last month of the quarter. In addition, nominal variables have been transformed to real terms using Shiller's CPI index. Annual dividends have been converted to quarterly figures by dividing them by 4 to compute quarterly returns. Data on trading volume is from Datastream. In particular, trading volume = Turnover by value (US-DS Market-TV) / Market value (US-DS Market - MV). Data on stock repurchases is from Refinitiv. Supply shocks are constructed using the repurchase rate $u_t \equiv \text{Repurchases}_t/\text{MarketCap}_t$ (both in dollars), which is unitless. Interpreting repurchases as proportionally retiring shares outstanding, I set the model's normalized supply to $\tilde{S}_t \equiv 1-u_t$ so that $\varepsilon_t^s = \tilde{S}_t - 1 = -u_t$. In turn, the discount factor has been calibrated to match the dynamics of the real T-Bill. The nominal risk-free rate is the 90-day T-Bill, from the FRED database and inflation expectations are from Michigan Surveys of Consumers, available here. I construct the ex-ante quarterly gross real T-bill return using the Fisher relation, combining the nominal gross return with expected gross inflation. The quarterly discount factor is then set to the inverse of this gross real return.

Macroeconomic data. Consumption data is the BEA real quarterly personal consumption expenditures series. Wages are the BEA compensation of employees.

Capital Tax Rates. The base effective average marginal rates on dividends, short and long capital gains and interests are supplied by the TAXSIM program of the National Bureau of Economic Research (NBER). See Feenberg and Coutts (1993) for a description of the program. The data can be found at here. These rates are offered on an annual basis from 1960 to 2018 at federal level and from 1979 to 2008 at state level. I took the rates computed using 1984 data in/deflated to law year, which reflects changes in tax rates without conflating changes in the income distribution.

Following Sialm (2009), I adjusted for state and local taxes before 1979 and after 2008 as well as for the distinction between qualified and nonqualified dividends from 2003 onward to get a complete series for the 1960-2018 period. Before 1960, the tax rates are interpolated using data from Sialm (2009). Details on the computation of the taxable share are available on Appendix F.

Realized Capital Gains. Total realized capital gains are from the Tax Foundation, which combines data from the Treasury Department, the Congressional Budget Office (CBO) and Wolters Kluwer. The series expands the years 1913-2025. The data uses returns with positive net capital gains and can be found here. The graph below compares this series with data from the CBO from here and here, and data directly from the IRS-

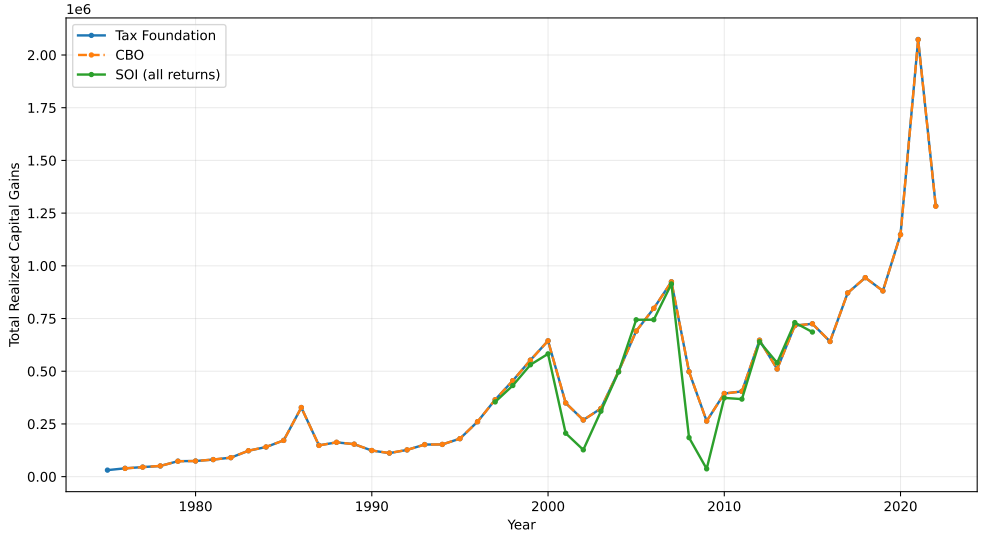


Figure 6: *Total Realized Capital Gains by source. Annual data.*

SOI (Table 1 in here) that covers the period 1997-2015. Crucially, the series from SOI includes all the returns (also the ones with net losses). As observed, all the series basically overlap, so using the series with net gains returns is without significant information loss.

Another concern is that the above series measure realized capital gains from all sources, whereas the model focuses on corporate equities. To assess whether this aggregation materially changes the relevant dynamics, I use SOI data that report realized gains by asset type. Over the sample, realized gains from corporate stock and mutual funds comove closely with total realized gains: the correlation in levels is 0.87, and the correlation in first differences is even higher at 0.96. The graph below illustrates this point. These results suggest that, although the levels differ, total realized gains are a good proxy for the time-series dynamics of equity realizations, making the macro elasticity estimated on total gains informative for the elasticity of corporate stock gains.

Under U.S. law, when net capital losses exceed net capital gains, taxpayers may use at most \$3,000 of the excess to reduce ordinary taxable income in the current year, carrying forward the remainder. In the model, I abstract from this contemporaneous ordinary-income offset while allowing loss carryovers, so the only omitted object is a bounded transfer of at most \$3,000 (times the marginal ordinary-income tax rate) in net-loss states. The table below shows that the capital-gains tax base is generated overwhelmingly by returns with very large net gains: AGI $\geq$ \$1 million filers account for about 73% of aggregate net gains, with average net gains per return on the order of \$1.5 million (roughly $500 \times 3{,}000$), and the extreme upper tail (AGI $\geq$ \$10 million) alone accounts for roughly 40% of aggregate net gains with average net gains around \$15 million (about $5{,}000 \times 3{,}000$). Consequently, the omitted \$3,000 offset is quantitatively immaterial at the scale that drives both aggregate revenues and the tax wedge in the model: the government meaningfully shares in upside

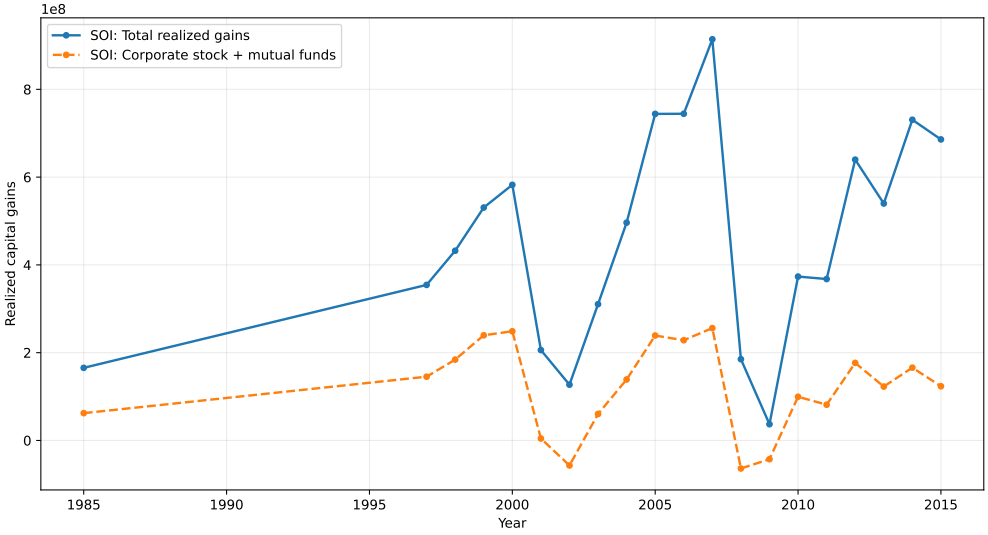


Figure 7: *Total vs Corporate Stock + Mutual Funds Realized Gains. Data from IRS-SOI.*

realizations through capital-gains taxation, but provides essentially no contemporaneous rebate in net-loss states. In this sense, for the economically relevant part of the distribution, the tax system is effectively asymmetric. To directly assess the size of the contemporaneous offset on net-loss returns, I use IRS aggregate net capital loss in AGI after limitation and the number of returns with net capital loss. Over 1998–2015, the average net loss that enters AGI per net-loss return is approximately \$2,223, with a range of roughly \$1,893–\$2,358. This is close to the statutory \$3,000 cap, indicating that net losses are overwhelmingly used to net against gains (or carried forward), while the contemporaneous ordinary-income offset is mechanically small.

Survey expectations. The main survey is the UBS Index of Investor Optimism. The quantitative question on stock market expectations has been surveyed over the period 1998:Q2-2007:Q4 with 702 responses per month on average. To make the data consistent with the model, I have run some adjustments. First, the series have been deflated by using inflation expectations from the Michigan Surveys of Consumers, available at https://data.sca.isr.umich.edu/data-archive/mine.php. Second, I transformed real returns expectations into capital gains expectations by subtracting an annualized dividend component. See Appendix G for more details.

A B. A Locked-In Bayesian Rational Expectations Equilibrium

Definition 2 (Locked-in B-REE)**.** *Fix an initial owner $j \in \{A, B\}$. A* lock-in *Bayesian–Rational Expectations equilibrium with owner j consists of:*

Table 8: *Net capital gains by AGI group: concentration and average net gain per return (SOI). The table is constructed from IRS SOI tabulations by AGI group. 'Avg net gain per return' equals total net gains in the group divided by the number of returns in the group. 'Share of returns' and 'share of net gains' are computed relative to totals across the AGI groups in the SOI extract for each year.*

Year	Share of returns	Share of net gains	Avg net gain per return ($)	Avg / $3,000
Panel A. AGI < $1 million				
2004	90.06%	28.46%	56,223	18.7×
2005	90.01%	28.85%	65,141	21.7×
2006	89.57%	28.63%	74,359	24.8×
2007	90.07%	23.70%	62,014	20.7×
Avg 2004–07	89.93%	27.41%	64,434	21.5×
Panel B. AGI ≥ $1 million				
2004	9.94%	71.54%	1,280,788	426.9×
2005	9.99%	71.15%	1,448,043	482.7×
2006	10.43%	71.37%	1,591,677	530.6×
2007	9.93%	76.30%	1,810,555	603.5×
Avg 2004–07	10.07%	72.59%	1,532,765	510.9×
Panel C. AGI ≥ $10 million				
2004	0.62%	35.67%	10,202,546	3,400.8×
2005	0.52%	34.30%	13,310,804	4,436.9×
2006	0.55%	39.43%	16,787,427	5,595.8×
2007	0.55%	48.93%	20,781,038	6,927.0×
Avg 2004–07	0.56%	39.58%	15,270,454	5,090.2×

1. *Asset holdings*

$$S_t^j = 1, \qquad S_t^{-j} = 0 \quad \textit{for all } t, \tag{102}$$

so the initial owner j holds the asset forever.

2. *State-contingent prices (P_H, P_L) such that, in each state $s \in \{H, L\}$,*

$$\textit{(no sale by } j\textit{):} \quad Q_s^j(\tau) \;\geq\; R_t^j(P_s;\tau) = (1-\tau)P_s + \tau P_{t-1}, \tag{103}$$

$$\textit{(no purchase by } -j\textit{):} \quad Q_s^{-j}(\tau) \;\leq\; P_s. \tag{104}$$

Combining (103) and (104), a pair (P_H, P_L) supports a locked-in B-REE with owner j if and only if, for each state $s \in \{H, L\}$,

$$Q_s^{-j}(\tau) \;\leq\; P_s \;\leq\; \frac{Q_s^j(\tau) - \tau P_{t-1}}{1-\tau}. \tag{105}$$

Whenever the interval in (105) is non-empty in both states,

$$Q_s^{-j}(\tau) \;\leq\; \frac{Q_s^j(\tau) - \tau P_{t-1}}{1-\tau} \qquad \text{for } s \in \{H, L\}, \tag{106}$$

there exists a continuum of locked-in B-REEs with owner j: any choice of

$$P_s \in \left[Q_s^{-j}(\tau),\ \left(Q_s^j(\tau) - \tau P_{t-1}\right)/(1-\tau) \right]$$

in each state yields an allocation with no trade and market-clearing asset demand.

In this regime, allocations are degenerate (the initial owner holds the asset forever) and prices are indeterminate within the no-trade bands (105); consequently, the level and volatility of prices are not uniquely determined.

C. Brief summary of changes in Capital Gains Taxation since 1975

Here I summarize all the tax reforms since 1975. In the mid-1970s, individuals excluded 50% of long-term gains; with a 70% top ordinary rate, the implied top rate on long-term gains was 35%. The Tax Reform Act of 1976 lengthened the long-term holding period from 6 to 9 months for 1977 and to 12 months beginning after 1978. The Revenue Act of 1978 increased the exclusion to 60% and repealed the add-on minimum tax on excluded gains, reducing the statutory top long-term rate to 28%.[52] The Economic Recovery Tax Act of 1981 (ERTA) further reduced the effective top rate by capping the maximum tax on net long-term gains at 20% for sales and exchanges after June 9, 1981. The Deficit Reduction Act of 1984 temporarily shortened the holding period back to 6 months for assets acquired after June 22, 1984 (through the end of 1987). The Tax Reform Act of 1986 eliminated the exclusion and taxed long-term gains as ordinary income; with a 28% top ordinary bracket, the effective long-term top rate was 28% starting in 1987. Subsequent budget acts in 1990 and 1993 raised top ordinary rates (to 31% and then 39.6%) but retained a 28% statutory maximum on net capital gains.

The Taxpayer Relief Act of 1997 lowered the top long-term rate to 20% (10% for taxpayers in the 15% ordinary bracket) for gains after May 6, 1997, but simultaneously increased the long-term holding-period threshold to 18 months for gains realized after July 28, 1997. However, the IRS Restructuring and Reform Act of 1998 restored the 12-month rule effective January 1, 1998.[53] From May 7, 1997 through 2002, the statutory top long-term rate for individuals was therefore 20%. EGTRRA 2001 introduced an 8% rate for five-year gains of taxpayers in the 15% bracket starting in 2001, and an 18% rate for

[52] As for the $3,000 limit: the first across-the-board limitation on using net capital losses to offset ordinary income was enacted in the Revenue Act of 1942, which set a $1,000 cap and created loss carryforwards; subsequent law kept the cap but later raised it. The Tax Reform Act of 1976 increased the cap to $2,000 for 1977 and to $3,000 for tax years beginning after 1977, which is the limit still in force for individuals today.

[53] In 1999 Congress passed the Taxpayer Refund and Relief Act of 1999 on August 5, 1999, which would have cut the long-term maximum to 18% (8% in the 15% bracket) and introduced limited indexation; President Clinton vetoed the bill on Sept. 23, 1999, leaving the 20%/10% schedule in place.

five-year gains on assets acquired after 2000 (which could not bind before 2006); these provisions did not alter the 20% top rate through 2002.

After EGTRRA, the Jobs and Growth Tax Relief Reconciliation Act of 2003 (JGTRRA) cut the statutory long-term capital-gains cap to 15% (and to 5% for taxpayers in the 10%–15% ordinary brackets), generally effective for gains realized after May 5, 2003; it also aligned qualified dividends with these same preferential rates. The five-year/18% regime scheduled under TRA97 (and the 8% companion for the low bracket) was effectively superseded by JGTRRA's unified 15%/5% structure. The Tax Increase Prevention and Reconciliation Act of 2005 (enacted May 17, 2006) extended the JGTRRA capital-gains/dividends rates through 2010 and confirmed the scheduled drop of the low-bracket rate to 0% beginning in 2008. The Tax Relief, Unemployment Insurance Reauthorization, and Job Creation Act of 2010 (TRA 2010) then extended the 0%/15% long-term structure through 2012. Beginning tax year 2013, the American Taxpayer Relief Act of 2012 (ATRA) restored a top ordinary bracket and set a new 20% top long-term rate for taxpayers above the ATRA thresholds, while retaining 0%/15% below those thresholds. Separately—but effective the same date— the Affordable Care Act's §1411 Net Investment Income Tax imposed an additional 3.8% surtax on net investment income (including realized capital gains) above statutory thresholds, leaving statutory rates at 0/15/20 but raising effective marginal burdens for high-income filers. From 2013 through 2022, the 0/15/20 framework and the one-year long-term holding-period remained intact.

D. Proofs of Results

A.1 Proof of Lemma 1

Proof. The constraint set Γ_t^i features a kink at $S_t^i = S_{t-1}^i$ due to the capital gains tax term, making it piecewise linear and non-differentiable. To handle this, I analyze the Bellman equation piecewise. Define

$$V^i(z_t^i) = \max_{0 \le S_t^i \le 1} \mathcal{W}(S_t^i; z_t^i), \tag{107}$$

where the return function is

$$\mathcal{W}(S_t^i; z_t^i) = \begin{cases} W_t + (D_t + P_t)S_{t-1}^i - P_t S_t^i + T_t^i + \delta \mathbb{E}^{\mathcal{P}^i}[V^i(z_{t+1}^i) \mid z_t^i] & \text{if } S_t^i \ge S_{t-1}^i \\ W_t + (D_t + P_t)S_{t-1}^i - P_t S_t^i - \tau(S_{t-1}^i - S_t^i)\max\{0, P_t - P_{t-1}\} & \\ \quad + T_t^i + \delta \mathbb{E}^{\mathcal{P}^i}[V^i(z_{t+1}^i) \mid z_t^i] & \text{if } S_t^i < S_{t-1}^i \end{cases}$$

A key object for characterizing optimal policies is the expected marginal value of shareholdings:

$$\mathbb{E}^{\mathcal{P}^i}\left[\frac{\partial V^i(z_{t+1}^i)}{\partial S_t^i} \,\middle|\, z_t^i\right].$$

While the standard envelope theorem does not apply due to the non-differentiability at $S^i_{t+1} = S^i_t$, the generalized envelope theorem of Milgrom and Segal (2002) yields

$$\frac{\partial V^i(z^i_t)}{\partial S^i_{t-1}} = \begin{cases} D_t + P_t & \text{if } S^i_t \geq S^i_{t-1} \\ D_t + P_t - \tau \max\{0, P_t - P_{t-1}\} & \text{if } S^i_t < S^i_{t-1} \end{cases} \tag{108}$$

Therefore, the expected marginal value satisfies

$$\mathbb{E}^{\mathcal{P}^i}\left[\frac{\partial V^i(z^i_{t+1})}{\partial S^i_t} \,\middle|\, z^i_t\right] = \mathbb{E}^{\mathcal{P}^i}\left[D_{t+1} + P_{t+1} - \tau \max\{0, P_{t+1} - P_t\} \mathbb{1}\{S^i_{t+1} < S^i_t\} \mid z^i_t\right]. \tag{109}$$

Define the investor's subjective valuation as

$$Q^i_t \equiv \delta \mathbb{E}^{\mathcal{P}^i}\left[\frac{\partial V^i(z^i_{t+1})}{\partial S^i_t} \,\middle|\, z^i_t\right], \tag{110}$$

which represents the discounted expected marginal value of holding an additional share.

The first-order conditions depend on the investor's initial position. For $S^i_{t-1} = 0$:

$$S^i_t = \begin{cases} \bar{S} & \text{if } Q^i_t > P_t \\ 0 & \text{if } Q^i_t < P_t \\ \{0, \bar{S}\} & \text{if } Q^i_t = P_t \end{cases} \tag{111}$$

For $S^i_{t-1} = \bar{S}$:

$$S^i_t = \begin{cases} \bar{S} & \text{if } Q^i_t > P_t - \tau \max\{0, P_t - P_{t-1}\} \\ 0 & \text{if } Q^i_t < P_t - \tau \max\{0, P_t - P_{t-1}\} \\ \{0, \bar{S}\} & \text{if } Q^i_t = P_t - \tau \max\{0, P_t - P_{t-1}\} \end{cases} \tag{112}$$

For interior initial holdings $S^i_{t-1} \in (0, \bar{S})$:

$$S^i_t = \begin{cases} \bar{S} & \text{if } Q^i_t > P_t \\ S^i_{t-1} & \text{if } P_t - \tau \max\{0, P_t - P_{t-1}\} \leq Q^i_t \leq P_t \\ 0 & \text{if } Q^i_t < P_t - \tau \max\{0, P_t - P_{t-1}\} \end{cases} \tag{113}$$

These conditions can be summarized in the stock policy function (7). ■

A.2 Proof of Proposition 1

Proof. The proof proceeds by construction. I first conjecture that the equilibrium features complete specialization, with type A agents holding the entire stock supply in state H and type B agents holding it in state L. Given this allocation pattern, I derive the unique

market-clearing prices by solving agents' Euler equations and then verify that, given these prices, the conjectured allocation pattern is indeed optimal for each agent type.

The parameter restriction $\pi^A_{H\to H} > 1-\pi^B_{L\to L}$ guarantees $P^*_H > P^*_L$, a natural order given $D_H > D_L$. This restriction along with the disagreement about persistence ($\pi^A_{H\to H} > \pi^B_{H\to H}$ and $\pi^A_{L\to L} > \pi^B_{L\to L}$) implies $Q^A_H > Q^B_H$ and $Q^B_L > Q^A_L$, where

$$Q^i_H = \delta[\pi^i_{H\to H}(D_H + P_H) + (1-\pi^i_{H\to H})(D_L + P_L)] \tag{114}$$

$$Q^i_L = \delta[\pi^i_{L\to L}(D_L + P_L) + (1-\pi^i_{L\to L})(D_H + P_H - \tau(P_H - P_L))] \tag{115}$$

for $i = A, B$. These conditions come from the subjective marginal valuation of each investor, noting two issues: i) $P_H > P_L$ implies taxes are only faced in the low-to-high transition; ii) common knowledge about beliefs, combined with rationality, implies that agents correctly anticipate time-invariant state prices given stationary fundamentals and beliefs, and then $\mathbb{E}^{\mathcal{P}^i}_t[P_{st+1}] = P_s$.

By market clearing,

$$P^*_H = Q^A_H \tag{116}$$

$$P^*_L = Q^B_L \tag{117}$$

Hence, equilibrium prices solve this system of equations:

$$P^*_H = \delta[\pi^A_{H\to H}(D_H + P^*_H) + (1-\pi^A_{H\to H})(D_L + P^*_L)] \tag{118}$$

$$P^*_L = \delta[\pi^B_{L\to L}(D_L + P^*_L) + (1-\pi^B_{L\to L})(D_H + P^*_H - \tau(P^*_H - P^*_L))] \tag{119}$$

The solution to this system is given by equation (14).

To see why equations (116)-(117) determines equilibrium prices, focus on $s_t = H$. In this case, if $P_H < Q^A_H$, there is excess demand, even if $P_H > Q^B_H$. At $P_H = Q^A_H$, agents in group A are willing to hold or buy any quantity, and since $Q^A_H > Q^B_H$, agents in group B are unwilling to buy or, in case they were the previous owners, are willing to sell, provided P_H is high enough to compensate for tax payments, a condition that will be verified below.

I proceed case-by-case to verify (P^*_H, P^*_L) delivers the conjectured allocations. First, consider state $s_t = H$ with $s_{t-1} = H$, such that $(S^A_{t-1}, S^B_{t-1}) = (1, 0)$. The basis price is $P_H = P^*_H$, which implies the no-trade region is $P^*_H - \tau(P^*_H - P^*_H) \leq Q^A_H \leq P^*_H$. Thus $Q^A_H = P^*_H$. Agent B remains out of the market if $Q^B_H \leq P^*_H$ which is strictly satisfied as $Q^B_H < Q^A_H$. Similarly, in $s_t = L$ with $s_{t-1} = L$ such that $(S^A_{t-1}, S^B_{t-1}) = (0, 1)$, by symmetric reasoning, $Q^B_L = P^*_L$ and type A optimally holds zero since $Q^A_L < Q^B_L$.

Then, consider $s_t = H$ with $s_{t-1} = L$ such that $(S^A_{t-1}, S^B_{t-1}) = (0, 1)$: Type A buys if $Q^A_H \geq P^*_H$, which holds with equality. Type B sells if prices are high enough to compensate

for tax payments: $Q_H^B \leq P_H^* - \tau(P_H^* - P_L^*)$. Note that

$$P_H^* - Q_H^B = \delta\big(\pi_{H\to H}^A - \pi_{H\to H}^B\big)\Big[(D_H - D_L) + (P_H^* - P_L^*)\Big]$$

Hence $Q_H^B \leq P_H^* - \tau(P_H^* - P_L^*)$ implies

$$\delta\left(\pi_{H\to H}^A - \pi_{H\to H}^B\right)\left[(D_H - D_L) + (P_H^* - P_L^*)\right] \geq \tau\left(P_H^* - P_L^*\right)$$

or equivalently

$$\pi_{H\to H}^A - \pi_{H\to H}^B \;\geq\; \frac{\tau}{\delta}\,\frac{P_H^* - P_L^*}{(D_H - D_L) + (P_H^* - P_L^*)} \tag{120}$$

The seller participation constraint therefore requires a minimum amount of disagreement between type A and type B beliefs: the belief wedge $\pi_{H\to H}^A - \pi_{H\to H}^B$ must be large enough to compensate the seller for the realization tax. Using the closed-form price gap, the right-hand side of (120) can be written in terms of beliefs. Define

$$\kappa\big(\tau, \pi_{H\to H}^A\big) \;\equiv\; \frac{\tau\Big(\pi_{H\to H}^A - \big(1 - \pi_{L\to L}^B\big)\Big)}{1 - \delta\,\tau\big(1 - \pi_{L\to L}^B\big)} \tag{121}$$

Then trading in the $L \to H$ transition requires

$$\pi_{H\to H}^A - \pi_{H\to H}^B \;\geq\; \kappa(\tau, \pi_{H\to H}^A) \tag{122}$$

Because $\kappa(\tau, \pi_{H\to H}^A)$ depends on $\pi_{H\to H}^A$, the wedge condition (122) can be expressed either as (i) a lower bound on $\pi_{H\to H}^A$ for a given $\pi_{H\to H}^B$, or (ii) an upper bound on $\pi_{H\to H}^B$ for a given $\pi_{H\to H}^A$. I follow ii) as it is slightly cleaner. To obtain a restriction that guarantees trade for all tax rates $\tau \in [0,1]$, note that under the maintained assumption $\pi_{H\to H}^A > 1 - \pi_{L\to L}^B$ and writing $\omega \equiv \pi_{H\to H}^A - (1 - \pi_{L\to L}^B) > 0$, one obtains

$$\kappa(\tau, \pi_{H\to H}^A) = \frac{\tau\,\omega}{1 - \delta\,\tau(1 - \pi_{L\to L}^B)} \;>\; 0 \quad \text{for } \tau > 0$$

and

$$\frac{\partial}{\partial\tau}\kappa(\tau, \pi_{H\to H}^A) = \frac{\omega}{\Big(1 - \delta\,\tau\big(1 - \pi_{L\to L}^B\big)\Big)^2} \;>\; 0 \tag{123}$$

Thus $\kappa(\tau, \pi_{H\to H}^A)$ is increasing in τ, so the tightest version of (122) over $\tau \in [0,1]$ occurs at $\tau = 1$. Evaluating (122) at $\tau = 1$ yields

$$\pi_{H\to H}^B \;\leq\; \pi_{H\to H}^A - \kappa(1, \pi_{H\to H}^A) \tag{124}$$

Using (121) at $\tau = 1$ gives

$$\kappa(1, \pi^A_{H\to H}) = \frac{\pi^A_{H\to H} - (1 - \pi^B_{L\to L})}{1 - \delta\big(1 - \pi^B_{L\to L}\big)} \tag{125}$$

Substituting (125) into (124) and rearranging yields the closed-form upper bound

$$\pi^B_{H\to H} \;\le\; \pi^A_{H\to H} - \kappa(1, \pi^A_{H\to H}) = \frac{\big(1 - \pi^B_{L\to L}\big)\big(1 - \delta\,\pi^A_{H\to H}\big)}{1 - \delta\big(1 - \pi^B_{L\to L}\big)} \tag{126}$$

The bound (126) is automatically feasible. The denominator is strictly positive for $\delta \in (0,1)$ and $\pi^B_{L\to L} \in [0,1]$, and the numerator is nonnegative since $1 - \pi^B_{L\to L} \ge 0$ and $1 - \delta\,\pi^A_{H\to H} \ge 0$ for $\pi^A_{H\to H} \le 1$. Hence (126) is compatible with $\pi^B_{H\to H} \ge 0$.

Finally, consider $s_t = L$ with $s_{t-1} = H$ such that $(S^A_{t-1}, S^B_{t-1}) = (1, 0)$: Type A sells (no tax since $P^*_L < P^*_H$) if $Q^A_L \le P^*_L$. Type B buys if $Q^B_L \ge P^*_L$. Both conditions are satisfied given $P^*_L = Q^B_L > Q^A_L$. ■

A.3 Proof of Proposition 2

Proof. In a stationary B-REE, the price takes only two values, $P^*_t \in \{P^*_H, P^*_L\}$, with $\text{Prob}(P^*_t = P^*_H) = \mu$. Hence

$$\text{Var}(P^*_t) = \mu(1-\mu)\Delta(\tau, \pi_A)^2 \tag{127}$$

where $\Delta(\tau, \pi_A) \equiv P^*_H(\tau, \pi_A) - P^*_L(\tau, \pi_A)$; the notation emphasizes the dependence on the relevant objects. The total effect of τ on it is given by:

$$\frac{d}{d\tau}\text{Var}(P^*_t) = 2\mu(1-\mu)\,\Delta(\tau, \pi_A(\tau))\,\frac{d\Delta(\tau, \pi_A(\tau))}{d\tau} \tag{128}$$

Since $\Delta(\tau, \pi_A) > 0$ in a B–REE (because $P^*_H > P^*_L$), it is sufficient to show $\frac{d\Delta}{d\tau} > 0$. Differentiating,

$$\frac{d\Delta}{d\tau} = \underbrace{\frac{\partial \Delta}{\partial \tau}}_{\text{holding beliefs fixed}} + \underbrace{\mathbf{1}\{\text{frontier binds}\}\,\frac{\partial \Delta}{\partial \pi_A}\,\pi'_A(\tau)}_{\text{belief adjustment along the binding frontier}} \tag{129}$$

where $\mathbf{1}\{\text{frontier binds}\} = 1$ only when the π_A frontier is binding, i.e. $\pi^B_{H\to H} + \kappa(\tau, \pi^A_{H\to H}) = \pi^A_{H\to H}$, and equals 0 otherwise.

Interior of the B-REE region. Start in the interior of the B-REE region, such that $\pi'_A(\tau) = 0$ and only the first component matters. Using $\mathbf{P}^* = K(\tau)\mathbf{D}^e$ from Proposition 1, rewrite equilibrium prices as

$$P^*_H(\tau, \pi_A) = \frac{\delta}{M(\tau)}\,A(\tau, \pi_A), \qquad P^*_L(\tau, \pi_A) = \frac{\delta}{M(\tau)}\,B(\tau, \pi_A),$$

where

$$A(\tau,\pi_A)\equiv\big(1-\delta[\pi_B+\tau(1-\pi_B)]\big)D_H^e+\delta(1-\pi_A)D_L^e,\qquad B(\tau,\pi_A)\equiv\delta(1-\pi_B)(1-\tau)D_H^e+(1-\delta\pi_A)D_L^e,$$

Hence

$$\Delta(\tau,\pi_A)\equiv P_H^*(\tau,\pi_A)-P_L^*(\tau,\pi_A)=\frac{\delta}{M(\tau)}\Big(A(\tau,\pi_A)-B(\tau,\pi_A)\Big).$$

Note the key simplification:

$$\begin{aligned}A(\tau,\pi_A)-B(\tau,\pi_A)&=\Big(1-\delta[\pi_B+\tau(1-\pi_B)]-\delta(1-\pi_B)(1-\tau)\Big)D_H^e\\&\quad+\Big(\delta(1-\pi_A)-(1-\delta\pi_A)\Big)D_L^e\\&=(1-\delta)D_H^e-(1-\delta)D_L^e\\&=(1-\delta)\big(D_H^e-D_L^e\big)\end{aligned}$$

so that

$$\Delta(\tau,\pi_A)=\frac{\delta(1-\delta)}{M(\tau)}\big(D_H^e-D_L^e\big)\tag{130}$$

Since $D_H^e-D_L^e=(\pi_A+\pi_B-1)(D_H-D_L)$, the maintained B-REE condition $\pi_A+\pi_B-1>0$ and $D_H>D_L$ imply $D_H^e>D_L^e$. Differentiating (130) w.r.t. τ holding beliefs fixed yields

$$\frac{\partial\Delta}{\partial\tau}=-\frac{\delta(1-\delta)\big(D_H^e-D_L^e\big)}{M(\tau)^2}M_\tau(\tau)$$

A direct differentiation of $M(\tau)$ in (9) gives

$$M_\tau(\tau)=-\delta(1-\pi_B)(1-\delta)$$

which is strictly negative under $\delta\in(0,1)$ and $\pi_B<1$. Hence $-M_\tau(\tau)>0$ and therefore

$$\frac{\partial\Delta}{\partial\tau}>0$$

Boundary-belief adjustment. Focus now at the boundary where $\pi_A=\pi_A(\tau)$. From (130), define

$$N(\pi_A)\equiv\delta(1-\delta)\big(D_H^e(\pi_A)-D_L^e\big)$$

such that $\Delta(\tau,\pi_A)=\frac{N(\pi_A)}{M(\tau,\pi_A)}$. Then,

$$\frac{\partial\Delta}{\partial\pi_A}=\frac{N_{\pi_A}(\pi_A)\,M(\tau,\pi_A)-N(\pi_A)\,M_{\pi_A}(\tau,\pi_A)}{M(\tau,\pi_A)^2}$$

Since $D_H^e(\pi_A)=D_L+\pi_A(D_H-D_L)$ and D_L^e is independent of π_A,

$$N_{\pi_A}(\pi_A)=\delta(1-\delta)(D_H-D_L)>0$$

Moreover, differentiating $M(\tau, \pi_A)$ from (9) w.r.t. π_A yields the simplification

$$\begin{aligned} M_{\pi_A}(\tau, \pi_A) &= (-\delta)\big(1 - \delta[\pi_B + \tau(1 - \pi_B)]\big) - \delta^2(-1)(1 - \pi_B)(1 - \tau) \\ &= -\delta(1 - \delta) \end{aligned}$$

which is strictly negative. Under the maintained restrictions, $M(\tau, \pi_A) > 0$ and (in the B-REE region) $N(\pi_A) > 0$, so both terms in the numerator are positive, implying

$$\frac{\partial \Delta}{\partial \pi_A} > 0$$

Finally, determine the sign of $\partial \pi_A / \partial \tau$. At the boundary, admissible beliefs satisfy $\pi^B_{H\to H} = \pi^A_{H\to H} - \kappa(\tau, \pi^A_{H\to H})$. Define the implicit function

$$F(\tau, \pi_A) \;\equiv\; \pi_A - \kappa(\tau, \pi_A) - \pi^B_{H\to H} \;=\; 0$$

which pins down $\pi_A = \pi_A(\tau)$ locally. By the implicit-function theorem,

$$\pi'_A(\tau) = -\frac{F_\tau(\tau, \pi_A)}{F_{\pi_A}(\tau, \pi_A)} = \frac{\kappa_\tau(\tau, \pi_A)}{1 - \kappa_{\pi_A}(\tau, \pi_A)} \tag{131}$$

It therefore suffices to show $\kappa_\tau(\tau, \pi_A) > 0$ and $1 - \kappa_{\pi_A}(\tau, \pi_A) > 0$. Given

$$\kappa(\tau, \pi_A) = \frac{\tau\left(\pi_A - (1 - \pi^B_{L\to L})\right)}{1 - \delta\,\tau\,(1 - \pi^B_{L\to L})}$$

one obtains

$$\kappa_\tau(\tau, \pi_A) = \frac{\pi_A - (1 - \pi^B_{L\to L})}{\Big(1 - \delta\,\tau\,(1 - \pi^B_{L\to L})\Big)^2} \;>\; 0 \tag{132}$$

and

$$\kappa_{\pi_A}(\tau, \pi_A) = \frac{\tau}{1 - \delta\,\tau\,(1 - \pi^B_{L\to L})} \;>\; 0 \tag{133}$$

where the signs follow from the restrictions on beliefs assumed in Proposition 1. The regularity condition $F_{\pi_A} = 1 - \kappa_{\pi_A}(\tau, \pi_A) > 0$ requires

$$1 - \frac{\tau}{1 - \delta\,\tau\,(1 - \pi^B_{L\to L})} > 0 \qquad \Longleftrightarrow \qquad \tau\left(1 + \delta(1 - \pi^B_{L\to L})\right) < 1 \tag{134}$$

which holds for sufficiently small τ and is maintained as a parameter restriction on the admissible belief boundary. By (131), together with (132) and (134), $\pi'_A(\tau) > 0$.

Since $\pi'_A(\tau) > 0$ on the boundary, the belief-adjustment term $\frac{\partial \Delta}{\partial \pi_A}\pi'_A(\tau)$ is strictly positive, and thus $\frac{d\Delta}{d\tau} > 0$ also on the boundary. Altogether, substituting $\frac{d\Delta}{d\tau} > 0$ into (128) yields $\frac{d}{d\tau}\mathrm{Var}(P^*_t) > 0$. ■

A.4 Proof of Lemma 2

Step 1: variance decomposition. Fix $s \in \{H, L\}$. Under dividend learning, equilibrium prices satisfy

$$\boldsymbol{P}_t^* = \frac{\delta}{M^D(\tau)} \mathcal{A}^D(\tau) \boldsymbol{\beta}_t^D$$

Let $a_s(\tau) \equiv e_s' \mathcal{A}^D(\tau)$ be the s-th pricing row and define

$$\bar{P}_{s,t}(\tau) \equiv a_s(\tau) \boldsymbol{\beta}_t^D \tag{135}$$

Then

$$P_{s,t}^* = \frac{\delta}{M^D(\tau)} \bar{P}_{s,t}(\tau) \tag{136}$$

Let $\Sigma_\beta^D \equiv \mathrm{Var}(\boldsymbol{\beta}_t^D)$ and $\bar{v}_s(\tau) \equiv \mathrm{Var}(\bar{P}_{s,t}(\tau))$. From (136),

$$\mathrm{Var}(P_{s,t}^*) = \left(\frac{\delta}{M^D(\tau)}\right)^2 \bar{v}_s(\tau) \tag{137}$$

Taking variance on both sides of (135) one obtains

$$\bar{v}_s(\tau) = a_s(\tau) \Sigma_\beta^D a_s(\tau)' \tag{138}$$

I first show that Σ_β^D is independent of τ. Consider the mean-coefficient approximation for dividend learning:

$$\beta_{s,t}^D = (1 - g\lambda_s)\beta_{s,t-1}^D + g\lambda_s \mathcal{D}_t$$

for $s \in \{H, L\}$. Let $\Lambda = \mathrm{diag}(\lambda_H, \lambda_L)$ and $A_\beta \equiv I_2 - g\Lambda$. Then

$$\boldsymbol{\beta}_t^D = A_\beta \boldsymbol{\beta}_{t-1}^D + g\Lambda \mathbf{1}\, \mathcal{D}_t$$

so Σ_β^D solves the Lyapunov equation

$$\Sigma_\beta^D = A_\beta \Sigma_\beta^D A_\beta' + g^2 \sigma_{\varepsilon^\mathcal{D}}^2\, \Lambda \mathbf{1}\mathbf{1}' \Lambda$$

hence (since $\lambda_H + \lambda_L = 1$)

$$\Sigma_\beta^D = \sigma_{\varepsilon^\mathcal{D}}^2 \begin{pmatrix} \dfrac{g\lambda_H}{2 - g\lambda_H} & \dfrac{g\lambda_H\lambda_L}{1 - g\lambda_H\lambda_L} \\ \dfrac{g\lambda_H\lambda_L}{1 - g\lambda_H\lambda_L} & \dfrac{g\lambda_L}{2 - g\lambda_L} \end{pmatrix}$$

which is independent of τ. Therefore,

$$\bar{v}_s'(\tau) = \frac{da_s(\tau)}{d\tau} \Sigma_\beta^D a_s(\tau)' + a_s(\tau) \Sigma_\beta^D \frac{da_s(\tau)'}{d\tau} = 2\, \frac{da_s(\tau)}{d\tau} \Sigma_\beta^D a_s(\tau)' \tag{139}$$

From the pricing coefficients, for both $s \in \{H, L\}$,

$$\frac{da_s(\tau)}{d\tau} = -\delta(1-\pi_B)\, q$$

with $q \equiv (\pi_A,\ 1-\pi_A)$. Hence,

$$\bar{v}_s'(\tau) = -2\delta(1-\pi_B)\, q\, \Sigma_\beta^D a_s(\tau)' \tag{140}$$

Under the maintained belief configuration in Proposition 1, the pricing loadings are strictly positive entrywise: $a_s(\tau) \gg 0$ for all $\tau \in [0,1]$ and $s \in \{H, L\}$. Checking that is immediate. The only potentially problematic term is

$$a_{HH}(\tau) = (1 - \delta\tau(1-\pi_B))\pi_A + \delta(1-\pi_A-\pi_B)$$

since $\pi_A + \pi_B > 1$ in Proposition 1 makes the last term negative. Because $a_{HH}(\tau)$ is decreasing in τ, it suffices to impose $a_{HH}(1) > 0$, which (when $\delta(2-\pi_B) > 1$) is equivalent to the upper bound

$$\pi_A \ < \ \frac{\delta(1-\pi_B)}{\delta(2-\pi_B)-1}$$

Since all entries of Σ_β^D are strictly positive, it follows that $\Sigma_\beta^D a_s(\tau)' > 0$ and, since $q \gg 0$, $q\Sigma_\beta^D a_s(\tau)' > 0$, implying

$$\bar{v}_s'(\tau) < 0 \quad \text{for all } \tau \in [0,1],\ s \in \{H, L\} \tag{141}$$

Differentiate (137):

$$\frac{d}{d\tau}\mathrm{Var}(P_{s,t}^*) = \Big(\frac{\delta}{M^D(\tau)}\Big)^2 \bar{v}_s'(\tau) + \bar{v}_s(\tau)\, \frac{d}{d\tau}\Big(\frac{\delta}{M^D(\tau)}\Big)^2 \tag{142}$$

From

$$M^D(\tau) = (1-\delta\pi_A)\big(1-\delta\tau(1-\pi_B)-\delta\pi_B\big) - \delta^2(1-\pi_A)(1-\pi_B)(1-\tau)$$

one obtains

$$(M^D)'(\tau) = -\delta(1-\pi_B)(1-\delta) < 0 \tag{143}$$

so $\delta/M^D(\tau)$ is increasing in τ and the second term in (142) is nonnegative. A sufficient condition for $\frac{d}{d\tau}\mathrm{Var}(P_{s,t}^*) \le 0$ is that the negative contribution from $\bar{v}_s'(\tau)$ dominates the positive rescaling term:

$$-\bar{v}_s'(\tau) \ \ge \ 2\Big(-\frac{(M^D)'(\tau)}{M^D(\tau)}\Big)\bar{v}_s(\tau) \tag{144}$$

Using (140), (143), and $\bar{v}_s(\tau) = a_s(\tau)\Sigma_\beta^D a_s(\tau)'$, the dominance condition (144) is equivalent to

$$N_s(\tau) \;\geq\; \frac{1-\delta}{M^D(\tau)}\, D_s(\tau) \tag{145}$$

where $N_s(\tau) \equiv q\,\Sigma_\beta^D a_s(\tau)'$ and $D_s(\tau) \equiv a_s(\tau)\Sigma_\beta^D a_s(\tau)'$. By construction, $q \geq 0$ and each entry of $a_s(\tau)$ is affine and decreasing in τ. Moreover, Σ_β^D has nonnegative entries (under $g\lambda_H < 2$, $g\lambda_L < 2$, and $g\lambda_H\lambda_L < 1$), and therefore both scalar maps $a \mapsto q\,\Sigma_\beta^D a'$ and $a \mapsto a\,\Sigma_\beta^D a'$ are increasing in each component of a on $\mathbb{R}^2_+$. Since $a_s(1) \leq a_s(\tau) \leq a_s(0)$ entrywise for all $\tau \in [0,1]$, it follows that, for all $\tau \in [0,1]$

$$\begin{aligned} N_s(\tau) \;&\geq\; N_s(1) = q\,\Sigma_\beta^D a_s(1)' \\ D_s(\tau) \;&\leq\; D_s(0) = a_s(0)\Sigma_\beta^D a_s(0)' \end{aligned} \tag{146}$$

Finally, since $(M^D)'(\tau) < 0$, we have $M^D(\tau) \geq M^D(1)$ and thus

$$\frac{1-\delta}{M^D(\tau)} \;\leq\; \frac{1-\delta}{M^D(1)} \qquad \forall\, \tau \in [0,1] \tag{147}$$

Combining (146)–(147), a sufficient condition for (145) to hold for all $\tau \in [0,1]$ is the pair of endpoint inequalities

$$q\,\Sigma_\beta^D a_s(1)' \;\geq\; \frac{1-\delta}{M^D(1)}\, a_s(0)\Sigma_\beta^D a_s(0)' \qquad \text{for each } s \in \{H, L\} \tag{148}$$

Equivalently, using $M^D(1) = (1-\delta\pi_A)(1-\delta)$,

$$F_s(\pi_A;\pi_B) \;\equiv\; (1-\delta\pi_A)\, q\,\Sigma_\beta^D a_s(1)' \;-\; a_s(0)\Sigma_\beta^D a_s(0)' \;\geq\; 0, \qquad s \in \{H, L\} \tag{149}$$

Fix π_B and let $[\underline{\pi}_A, \bar{\pi}_A]$ denote the admissible belief interval from Proposition 1. For each $s \in \{H, L\}$, define the cubic polynomial $F_s(\pi_A;\pi_B)$ by (149) and let

$$\mathcal{I}^D(\pi_B) \;\equiv\; \Big\{\pi_A \in [\underline{\pi}_A, \bar{\pi}_A] :\; F_H(\pi_A;\pi_B) \geq 0 \text{ and } F_L(\pi_A;\pi_B) \geq 0\Big\}$$

Because $F_H(\cdot;\pi_B)$ and $F_L(\cdot;\pi_B)$ are continuous in π_A, the set $\mathcal{I}^D(\pi_B)$ is closed (as an intersection of closed sets).[54] Since F_s is cubic in π_A, $\mathcal{I}^D(\pi_B)$ need not be connected in general. Accordingly, I impose the sufficient restriction directly as

$$\pi_A \;\in\; \mathcal{I}^D(\pi_B) \tag{151}$$

[54] In particular, $\mathcal{I}^D(\pi_B) \neq \emptyset$ whenever the endpoint $\pi_A = \bar{\pi}_A$ satisfies

$$F_H(\bar{\pi}_A;\pi_B) \geq 0 \qquad \text{and} \qquad F_L(\bar{\pi}_A;\pi_B) \geq 0. \tag{150}$$

Under (150), $\bar{\pi}_A \in \mathcal{I}^D(\pi_B)$.

Under (151), the inequalities in (149) hold for both $s \in \{H, L\}$, hence (148) holds. Therefore, (145) holds for all $\tau \in [0,1]$ and both $s \in \{H, L\}$, which together with (142) implies $\frac{d}{d\tau}\mathrm{Var}(P^*_{s,t}) \leq 0$ on $[0,1]$. The inequality is strict if (149) holds strictly for both $s \in \{H, L\}$ (equivalently, if π_A lies in the interior of the feasible set in (**??**)). ■

A.5 Proof of Lemma 3

Proof. *Proof.* For each state s,

$$\mathrm{Var}(P_{st}) = \delta^2 \, \mathcal{B}_s \, \Sigma_\beta \, \mathcal{B}_s' \tag{152}$$

Differentiating (202) w.r.t. τ,

$$\begin{aligned} \frac{d}{d\tau}\mathrm{Var}(P_{st}) &= \delta^2 \Big\{ \Big(\frac{\partial \mathcal{B}_s}{\partial \tau}\Big) \Sigma_\beta \, \mathcal{B}_s' + \mathcal{B}_s \, \Sigma_\beta \Big(\frac{\partial \mathcal{B}_s}{\partial \tau}\Big)' \; + \; \mathcal{B}_s \, \frac{d\Sigma_\beta}{d\tau} \, \mathcal{B}_s' \Big\} \\ &= 2\delta^2 \Big(\frac{\partial \mathcal{B}_s}{\partial \tau}\Big) \Sigma_\beta \, \mathcal{B}_s' \; + \; \delta^2 \, \mathcal{B}_s \, \frac{d\Sigma_\beta}{d\tau} \, \mathcal{B}_s' \end{aligned} \tag{153}$$

The first term captures the loadings effect (how τ changes $\mathcal{B}_s$); the second term captures how τ changes the beliefs' variance Σ_β (through learning dynamics).

Effect of τ on loadings

From (30), the H-row is $\mathcal{B}_H = (\pi_A, \; 1 - \pi_A)$, which does not depend on τ. In turn, the L-row elements $\mathcal{B}_L = (b_1, b_2)$ read as

$$b_1(\tau) = \frac{a(1-\tau)}{1 - \delta a \tau}; \qquad b_2(\tau) = \frac{\pi_B}{1 - \delta a \tau} \tag{154}$$

where $a \equiv 1 - \pi_B$. Let $\Sigma_\beta = \begin{pmatrix} \sigma_{HH} & \sigma_{HL} \\ \sigma_{HL} & \sigma_{LL} \end{pmatrix}$ and define

$$v \equiv \Sigma_\beta \, \mathcal{B}_L' = \begin{pmatrix} \sigma_{HH} b_1 + \sigma_{HL} b_2 \\ \sigma_{HL} b_1 + \sigma_{LL} b_2 \end{pmatrix} = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \tag{155}$$

Then

$$\frac{d}{d\tau}\mathrm{Var}(P_{Lt})\Big|_{\mathcal{B}} = 2\delta^2 \Big(\frac{\partial b_1}{\partial \tau} v_1 + \frac{\partial b_2}{\partial \tau} v_2 \Big) \tag{156}$$

A direct calculation yields

$$\frac{\partial b_1}{\partial \tau} = \frac{a(\delta a - 1)}{(1 - \delta a \tau)^2}; \qquad \frac{\partial b_2}{\partial \tau} = \frac{\delta a \, \pi_B}{(1 - \delta a \tau)^2} \tag{157}$$

Under $\delta a < 1$ (true generically since $\delta < 1$ and $a \leq 1$), we have $\partial b_1/\partial\tau < 0$ and $\partial b_2/\partial\tau > 0$. Expanding (156),

$$\frac{1}{2\delta^2}\frac{d}{d\tau}\mathrm{Var}(P_{Lt})\Big|_{\mathcal{B}} = \underbrace{\Big(\frac{\partial b_1}{\partial\tau}b_1\Big)\sigma_{HH}}_{\text{(i) own-}H} + \underbrace{\Big(\frac{\partial b_2}{\partial\tau}b_2\Big)\sigma_{LL}}_{\text{(ii) own-}L} + \underbrace{\Big(\frac{\partial b_1}{\partial\tau}b_2 + \frac{\partial b_2}{\partial\tau}b_1\Big)\sigma_{HL}}_{\text{(iii) covariance}} \tag{158}$$

I now show that the overall sign of (158) is negative under the stated assumptions, that is, $\frac{d}{d\tau}\mathrm{Var}(P_{Lt})\big|_{\mathcal{B}} < 0$.

Covariance contribution. Provided $\sigma_{HL} \geq 0$ (shown below), the covariance contribution (iii) is negative iff

$$\frac{\partial b_1}{\partial\tau}\,b_2 + \frac{\partial b_2}{\partial\tau}\,b_1 < 0 \tag{159}$$

Using (154)–(157) and canceling the positive factor $(1-\delta a\tau)^{-3}$, this condition reduces to

$$\frac{\partial b_1}{\partial\tau}\,b_2 + \frac{\partial b_2}{\partial\tau}\,b_1 < 0 \iff 1 > \delta(1-\pi_B)\,(2-\tau) \tag{160}$$

Thus, the following lower bound on π_B ensures a negative covariance contribution:

$$\pi_B \;>\; \underline{\pi}_B \equiv 1 - \frac{1}{\delta(2-\tau)} \tag{161}$$

If $\delta(2-\tau) \leq 1$, the bound is nonbinding and the covariance term is negative for all $\pi_B \in [0,1)$.

Own-variance contributions. Write the sum of the two own-variance terms in (158) as

$$\Big(\frac{\partial b_1}{\partial\tau}b_1\Big)\sigma_{HH} \;+\; \Big(\frac{\partial b_2}{\partial\tau}b_2\Big)\sigma_{LL}$$

Since $\frac{\partial b_1}{\partial\tau} < 0 < \frac{\partial b_2}{\partial\tau}$, a sufficient condition for the own-variance contribution to be negative is

$$\Big|\frac{\partial b_1}{\partial\tau}\,b_1\Big|\,\sigma_{HH} \;>\; \Big(\frac{\partial b_2}{\partial\tau}\,b_2\Big)\,\sigma_{LL} \tag{162}$$

Define the variance ratio

$$\rho \;\equiv\; \frac{\sigma_{LL}}{\sigma_{HH}} \;\geq\; 0$$

Dividing (162) by σ_{HH} yields the equivalent condition

$$\Big|\frac{\partial b_1}{\partial\tau}\,b_1\Big| \;>\; \rho\Big(\frac{\partial b_2}{\partial\tau}\,b_2\Big) \tag{163}$$

Using (154)–(157) and canceling the common positive factor $(1-\delta a\tau)^{-3}$, condition (163) reduces to

$$-\delta\big((1-\tau)+\rho\big)\,\pi_B^2 \;+\; (1-\tau)(2\delta-1)\,\pi_B \;+\; (1-\tau)(1-\delta) \;>\; 0 \tag{164}$$

This inequality holds for π_B below its positive root,

$$\pi_B \;<\; \bar{\pi}_B(\rho) \;\equiv\; \frac{(1-\tau)(2\delta-1)+\sqrt{(1-\tau)^2(2\delta-1)^2+4\delta\big((1-\tau)+\rho\big)(1-\tau)(1-\delta)}}{2\delta\big((1-\tau)+\rho\big)} \tag{165}$$

Hence, the condition

$$\underline{\pi}_B \;<\; \pi_B \;<\; \bar{\pi}_B(\rho)$$

is sufficient to ensure that τ has a negative effect on the loadings through the own-variance terms. For $\tau \in (0,1)$ and $\delta < 1$, the belief interval $[\underline{\pi}_B, \bar{\pi}_B(\rho)]$ is non-empty if and only if $\bar{\pi}_B(\rho) \geq \underline{\pi}_B$.

Effect of τ on the beliefs covariance matrix

I first replace time-varying updating intensities by their stationary means so that beliefs follow a time-invariant linear recursion with a well-defined unconditional covariance. I then differentiate the associated covariance fixed point and study its sign in the small-gain region, where the discrete-time object admits a continuous-time approximation. Finally, I use a positivity argument to translate the sign condition into simple inequalities on primitives.

Mean–coefficient approximation and a Lyapunov representation. I use a mean–coefficient approximation that replaces the time-varying intensity matrix $\mathrm{diag}(\mathbf{I}_{t-1})$ by its stationary mean

$$\Lambda \;\equiv\; \mathbb{E}[\mathrm{diag}(\mathbf{I}_{t-1})] \;=\; \mathrm{diag}(\lambda_H, \lambda_L)$$

Under this approximation, beliefs follow the time-invariant VAR(1)

$$\beta_t \;=\; d \;+\; U(\tau)\beta_{t-1} \;+\; \eta_t,$$

with $U(\tau) = I_2 - G\Lambda + \delta G\Lambda\mathcal{B}(\tau)$, $\mathbb{E}[\eta_t] = 0$ and $\mathrm{Var}(\eta_t) = \Sigma_u = \sigma_b^2\, G\Lambda G$. Given the spectral radius of the matrix $U(\tau)$ is below 1,[55] the unconditional covariance $\Sigma_\beta \equiv \mathrm{Var}(\beta_t)$

[55] Since $U(\tau) = I_2 - G\Lambda + \delta G\Lambda\mathcal{B}(\tau)$, the i-th row sum satisfies

$$\sum_j U_{ij}(\tau) = 1 - g_i\lambda_i + \delta g_i\lambda_i \sum_j \mathcal{B}_{ij}(\tau) = 1 - g_i\lambda_i\Big(1 - \delta\sum_j \mathcal{B}_{ij}(\tau)\Big)$$

exists and solves the discrete Lyapunov equation

$$\Sigma_\beta \;=\; U(\tau)\,\Sigma_\beta\,U(\tau)^\top \;+\; \Sigma_u \tag{166}$$

Differentiating with respect to τ. Differentiating (166) with respect to τ (note that Σ_u is τ-invariant under the mean–coefficient approximation) gives

$$\frac{d\Sigma_\beta}{d\tau} \;=\; U\,\frac{d\Sigma_\beta}{d\tau}\,U^\top \;+\; W \tag{167}$$

with $W \equiv U_\tau \Sigma_\beta U^\top + U\Sigma_\beta U_\tau^\top$ and $U_\tau \equiv \frac{dU}{d\tau}$. Using the same Neumann series representation,

$$\frac{d\Sigma_\beta}{d\tau} \;=\; \sum_{k\geq 0} U^k\,W\,\big(U^\top\big)^k \tag{168}$$

This representation makes clear that the sign of $d\Sigma_\beta/d\tau$ is governed by the term W.

Small–gain limit and continuous-time approximation. To obtain a transparent characterization in the small-gain region, impose symmetric gains $g_H = g_L \equiv g > 0$ so that $G = gI_2$, and consider $g \to 0$. Let $Q \equiv \sigma_b^2 \Lambda$ and $A(\tau) \equiv \Lambda\big(I_2 - \delta\mathcal{B}(\tau)\big)$. Then $\Sigma_u = g^2 Q$ and $U(\tau) = I_2 - gA(\tau)$. $A(\tau)$ is nonsingular with positive diagonal and nonpositive off-diagonal entries.[56]

Then, the discrete-time covariance admits the scaling

$$\Sigma_\beta(\tau, g) = g\,\bar{\Sigma}(\tau) + O(g^2) \tag{169}$$

as $g \to 0$, where $\bar{\Sigma}(\tau)$ is the unique symmetric solution to the continuous-time Lyapunov equation

$$A(\tau)\,\bar{\Sigma}(\tau) + \bar{\Sigma}(\tau)\,A(\tau)^\top \;=\; Q \tag{170}$$

Moreover,

$$\frac{d\Sigma_\beta}{d\tau}(\tau, g) = g\,\bar{\Sigma}_\tau(\tau) + O(g^2) \tag{171}$$

Because $\sum_j \mathcal{B}_{1j}(\tau) = 1$ and $\sum_j \mathcal{B}_{2j}(\tau) < 1$ for $\tau > 0$, it follows that

$$\sum_j U_{1j}(\tau) = 1 - g_H\lambda_H(1-\delta) < 1, \qquad \sum_j U_{2j}(\tau) < 1$$

whenever $g_i\lambda_i > 0$ and $\delta \in (0,1)$. In the small-gain region (in particular, for g sufficiently small), $U(\tau)$ is entrywise nonnegative, so $\|U(\tau)\|_\infty = \max_i \sum_j U_{ij}(\tau) < 1$. Therefore $\rho(U(\tau)) \leq \|U(\tau)\|_\infty < 1$.

[56] For $\tau > 0$, $\mathcal{B}(\tau)$ is entrywise nonnegative and substochastic: its first row sums to one, while its second-row sum is

$$\sum_j \mathcal{B}_{2j}(\tau) = \frac{1-(1-\pi_B)\tau}{1-\delta(1-\pi_B)\tau} < 1 \qquad (\delta \in (0,1))$$

Moreover, for $\tau \in (0,1)$ we have $\mathcal{B}_{12}(\tau) > 0$ and $\mathcal{B}_{21}(\tau) > 0$, so the strict mass loss in the second row is reachable from both states, implying $\rho(\mathcal{B}(\tau)) < 1$. Hence $I_2 - \delta\mathcal{B}(\tau)$ is nonsingular. Its diagonal entries satisfy $1 - \delta\mathcal{B}_{ii}(\tau) > 0$ and its off-diagonal entries equal $-\delta\mathcal{B}_{ij}(\tau) \leq 0$, so $A(\tau) = \Lambda(I_2 - \delta\mathcal{B}(\tau))$ is nonsingular with positive diagonal and nonpositive off-diagonal entries.

with $\bar{\Sigma}_\tau(\tau) \equiv \frac{d\bar{\Sigma}(\tau)}{d\tau}$. Differentiating (170) yields

$$A\,\bar{\Sigma}_\tau + \bar{\Sigma}_\tau A^\top \;=\; -\Big(A_\tau \bar{\Sigma} + \bar{\Sigma} A_\tau^\top\Big) \tag{172}$$

with $A_\tau(\tau) \equiv \frac{dA(\tau)}{d\tau} = -\delta\,\Lambda\,\mathcal{B}_\tau(\tau)$. Thus, in the small-gain region the comparative statics of Σ_β with respect to τ are pinned down by the continuous-time resolvent associated with the linear operator $X \mapsto AX + XA^\top$ and the forcing matrix $A_\tau\bar{\Sigma} + \bar{\Sigma}A_\tau^\top$.

Positivity of the resolvent and a sufficient condition for $\bar{\Sigma}_\tau \le 0$. Because $A(\tau)$ has strictly positive diagonal entries, nonpositive off–diagonals, and is nonsingular with $A(\tau)^{-1} \ge 0$ entrywise, its exponential satisfies $e^{-A(\tau)s} \ge 0$ entrywise for all $s \ge 0$. Hence (170) admits the integral representation

$$\bar{\Sigma}(\tau) = \int_0^\infty e^{-A(\tau)s}\,Q\,e^{-A(\tau)^\top s}\,ds \tag{173}$$

which implies $\bar{\Sigma}(\tau) \ge 0$ entrywise. Similarly, (172) admits

$$\bar{\Sigma}_\tau(\tau) = -\int_0^\infty e^{-A(\tau)s}\left(A_\tau(\tau)\bar{\Sigma}(\tau) + \bar{\Sigma}(\tau)A_\tau(\tau)^\top\right)e^{-A(\tau)^\top s}\,ds \tag{174}$$

Therefore, a convenient sufficient condition for $\bar{\Sigma}_\tau(\tau) \le 0$ entrywise is

$$\Big(A_\tau\bar{\Sigma} + \bar{\Sigma}A_\tau^\top\Big) \ge 0 \text{ entrywise} \quad\Longrightarrow\quad \bar{\Sigma}_\tau \le 0 \text{ entrywise} \tag{175}$$

Combined with (171), this delivers $\frac{d\Sigma_\beta}{d\tau}(\tau, g) \le 0$ entrywise for all sufficiently small g.

Forcing term and primitive inequalities. Write

$$\bar{\Sigma}(\tau) = \begin{pmatrix} x & y \\ y & z \end{pmatrix}$$

with $x > 0,\ z > 0,\ y \ge 0$. Remember the second row of $\mathcal{B}(\tau)$ is $(b_1(\tau), b_2(\tau))$, where

$$b_1(\tau) = \frac{(1-\pi_B)(1-\tau)}{1-\delta(1-\pi_B)\tau}; \qquad b_2(\tau) = \frac{\pi_B}{1-\delta(1-\pi_B)\tau}$$

A direct differentiation gives $b_1'(\tau) < 0 < b_2'(\tau)$ whenever $\delta(1-\pi_B) < 1$. Define

$$\alpha(\tau) \equiv -\delta\lambda_L\,b_1'(\tau) > 0; \qquad \beta(\tau) \equiv \delta\lambda_L\,b_2'(\tau) > 0$$

Since $\mathcal{B}_\tau$ has only the second row nonzero, $A_\tau = -\delta\Lambda\mathcal{B}_\tau$ takes the form

$$A_\tau(\tau) = \begin{pmatrix} 0 & 0 \\ \alpha(\tau) & -\beta(\tau) \end{pmatrix}$$

It follows that

$$A_\tau\bar{\Sigma} + \bar{\Sigma}A_\tau^\top = \begin{pmatrix} 0 & \alpha x - \beta y \\ \alpha x - \beta y & 2(\alpha y - \beta z) \end{pmatrix} \tag{176}$$

Hence, the forcing matrix in (176) is entrywise nonnegative if and only if its two distinct off-diagonal and lower-right entries are nonnegative, i.e.

$$\alpha(\tau)\, x - \beta(\tau)\, y \geq 0 \qquad \text{and} \qquad \alpha(\tau)\, y - \beta(\tau)\, z \geq 0 \tag{177}$$

Since $\alpha(\tau), \beta(\tau) > 0$ and $\bar{\Sigma}$ is entrywise nonnegative with $x > 0$ and $y > 0$, these inequalities are equivalently expressed in ratio form as

$$\frac{\alpha(\tau)}{\beta(\tau)} \geq \frac{y}{x} \tag{178a}$$

$$\frac{\alpha(\tau)}{\beta(\tau)} \geq \frac{z}{y} \tag{178b}$$

Moreover, the loading ratio α/β is primitive and τ-invariant:

$$\frac{\alpha(\tau)}{\beta(\tau)} = \frac{-b_1'(\tau)}{b_2'(\tau)} = \frac{1 - \delta(1 - \pi_B)}{\delta\pi_B} \tag{179}$$

Thresholds in terms of $\kappa = \lambda_H/\lambda_L$. Let $\kappa \equiv \lambda_H/\lambda_L > 0$ and write

$$A(\tau) = \Lambda\big(I_2 - \delta\mathcal{B}(\tau)\big) = \begin{pmatrix} \lambda_H(1 - \delta\pi_A) & -\lambda_H\delta(1 - \pi_A) \\ -\lambda_L\delta b_1(\tau) & \lambda_L(1 - \delta b_2(\tau)) \end{pmatrix} \equiv \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

Since $Q = \sigma_b^2\,\mathrm{diag}(\lambda_H, \lambda_L)$, (170) is equivalent to the 3×3 linear system

$$a_{11}x + a_{12}y = \frac{\sigma_b^2\lambda_H}{2} \tag{180a}$$

$$(a_{11} + a_{22})y + a_{12}z + a_{21}x = 0 \tag{180b}$$

$$a_{21}y + a_{22}z = \frac{\sigma_b^2\lambda_L}{2} \tag{180c}$$

Solving (180) in closed form yields (x, y, z) and thus the ratios entering (177) as linear-fractional functions of κ. Substituting those expressions together with (179) into (177) and solving for κ delivers primitive thresholds Θ_1 and Θ_2:

$$\Theta_1 = \frac{1 - \delta\big(\pi_B + \tau(1-\pi_B)\big)}{\Xi} \tag{181}$$

and

$$\Theta_2 \;=\; \frac{\delta\,\pi_B\Big[\pi_A - (1-\tau)(1-\pi_B)\Big] + \Big(1 - 2\pi_B\Big) - \tau(1-\pi_B)}{(1-\delta\pi_A)(\pi_A + \pi_B - 1)\big(1 - \delta(1-\pi_B)\tau\big)}$$

Here Ξ is the (nonzero) coefficient on κ that arises in the first inequality in (178) after substituting the closed-form solution of (180); explicitly,

$$\begin{aligned}\Xi = {} & 1 - \delta\pi_A + \delta^2(\pi_A - \pi_B)(\pi_A + \pi_B - 1) \\ & + \tau\,\delta(\pi_B - 1)\Big[1 + \delta(\pi_B - 1) + \delta^2(\pi_A - 1)(\pi_A + \pi_B - 1)\Big]\end{aligned} \tag{182}$$

Defining

$$\Theta \equiv \max\{\Theta_1, \Theta_2\}$$

one obtains the sufficient condition

$$\kappa \geq \Theta \quad \Longrightarrow \quad A_\tau \bar{\Sigma} + \bar{\Sigma} A_\tau^\top \geq 0 \text{ entrywise} \quad \Longrightarrow \quad \bar{\Sigma}_\tau \leq 0 \text{ entrywise.}$$

By (171), it follows that for g sufficiently small,

$$\frac{d\Sigma_\beta}{d\tau}(\tau, g) \leq 0 \text{ entrywise}$$

■

A.6 Proof of Proposition 4

Proof. Plugging the equilibrium pricing function (78) into the beliefs updating equation (77) and rearranging

$$\beta_t = \frac{\delta}{1-\delta\tau} g D^e + \Big[1 - g\Big(1 - \frac{\delta(1-\tau)}{1-\delta\tau}\Big)\Big]\beta_{t-1} + g\varepsilon_t^b \tag{183}$$

This is a 1st order difference equation that fully characterizes the equilibrium dynamics of price beliefs. Stability requires:

$$-1 < 1 - g\Big(1 - \frac{\delta(1-\tau)}{1-\delta\tau}\Big) < 1$$

or, rearranging the terms,

$$0 < g\Big(1 - \frac{\delta(1-\tau)}{1-\delta\tau}\Big) < 2$$

Given the parameter domain $g \in (0,1)$, $\delta \in (0,1)$ and $\tau \in (0,1)$, the term in brackets is guaranteed to be in the $(0,1)$ interval, so that the inequalities are automatically satisfied.

Since the difference equation is stable, the variance is well defined. The closed-form reads as

$$\mathrm{Var}(\beta_t) = \frac{g}{\frac{1-\delta}{1-\delta\tau}\left(2 - g\frac{1-\delta}{1-\delta\tau}\right)}\sigma_b^2 \tag{184}$$

Hence,

$$\mathrm{Var}(P_t) = \omega\mathrm{Var}(\beta_t) = \frac{\delta^2(1-\tau)^2\, g\, \sigma_b^2}{(1-\delta)(1-\delta\tau)\left(2 - g\,\dfrac{1-\delta}{1-\delta\tau}\right)} \tag{185}$$

The derivatives with respect to τ are given by:

$$\frac{d}{d\tau}\omega = 2\Big(\frac{\delta(1-\tau)}{1-\delta\tau}\Big)\frac{\delta(\delta-1)}{(1-\delta\tau)^2} < 0$$

given $\delta < 1$. In turn,

$$\frac{d}{d\tau}\mathrm{Var}(\beta_t) = -\frac{2\,g\,\sigma_b^2\,\delta}{1-\delta}\,\frac{1 - g\,\dfrac{1-\delta}{1-\delta\tau}}{\left(2 - g\,\dfrac{1-\delta}{1-\delta\tau}\right)^2} < 0$$

Both derivatives imply $\frac{d}{d\tau}\mathrm{Var}(P_t) < 0$.

■

A.7 Proof of Proposition 5

Proof. The proof proceeds by construction, following the same conjecture-solve-verify steps as in the proposition for CGT.

Step 1: Conjecture full specialization and compute subjective valuations. Conjecture that the equilibrium allocation and price rule are as in (88). Under common knowledge of beliefs and rationality, investors correctly anticipate that $P_{t+1} = P_H$ if $s_{t+1} = H$ and $P_{t+1} = P_L$ if $s_{t+1} = L$.

Under the conjecture, portfolios are fully specialized: in state H type A holds the single share and type B holds none, while in state L the opposite holds. Hence, whether an agent trades next period is pinned down by the next-period state. In particular, conditional on today's state s_t: (i) the current owner sells next period only if the state switches, and (ii) the current non-owner buys next period only if the state switches. Therefore, a marginal share held at the end of period t delivers next period the dividend plus the continuation value, where the continuation value depends on whether the marginal share is carried into the next state (no trade) or liquidated at the after-tax resale price (trade). Using (83), the relevant subjective marginal valuations are:

Type A in H. A marginal share held in H is carried into the portfolio if $H \to H$ (no trade). If $H \to L$, type A becomes the seller next period and liquidates the marginal share

at the after-tax resale price $(1-\tau^F)P_L$. Thus

$$Q_H^A = \delta\Big[\pi_{H\to H}^A(D_H + P_H) + (1-\pi_{H\to H}^A)(D_L + (1-\tau^F)P_L)\Big] \tag{186}$$

Type B in L. Symmetrically, a marginal share held in L is carried into the portfolio if $L \to L$ (no trade), while if $L \to H$ type B becomes the seller next period and liquidates at $(1-\tau^F)P_H$. Hence

$$Q_L^B = \delta\Big[\pi_{L\to L}^B(D_L + P_L) + (1-\pi_{L\to L}^B)(D_H + (1-\tau^F)P_H)\Big] \tag{187}$$

Type B in H. Type B is the non-owner in H under the conjecture. If $H \to H$, the share is liquidated next period at the after-tax resale price $(1-\tau^F)P_H$ (since under specialization type B sells in H). If $H \to L$, the agent becomes the buyer next period; holding the share today avoids paying the taxed purchase price $(1+\tau^F)P_L$ tomorrow. Thus

$$Q_H^B = \delta\Big[\pi_{H\to H}^B(D_H + (1-\tau^F)P_H) + (1-\pi_{H\to H}^B)(D_L + (1+\tau^F)P_L)\Big] \tag{188}$$

Type A in L. Analogously, type A is the non-owner in L. If $L \to L$, holding an extra share implies liquidation next period at $(1-\tau^F)P_L$ under specialization; if $L \to H$, holding the share avoids a taxed purchase at $(1+\tau^F)P_H$. Therefore

$$Q_L^A = \delta\Big[\pi_{L\to L}^A(D_L + (1-\tau^F)P_L) + (1-\pi_{L\to L}^A)(D_H + (1+\tau^F)P_H)\Big] \tag{189}$$

Step 2: Solve for market-clearing prices. Consider state H on a switch $L \to H$: then type A enters with $S_{t-1}^A = 0$ and must buy 1. By the individual demand rule (82), buying requires $Q_H^A \geq (1+\tau^F)P_H$. Market clearing with full specialization pins down the marginal price at equality:

$$(1+\tau^F)P_H = Q_H^A \tag{190}$$

Likewise, in state L on a switch $H \to L$, type B enters with $S_{t-1}^B = 0$ and must buy 1, so

$$(1+\tau^F)P_L = Q_L^B \tag{191}$$

These two equations determine the stationary price vector because the state-contingent price rule depends only on s_t; on non-switch histories there is no trade under full specialization, but the same (P_H, P_L) must prevail for investors to correctly anticipate tomorrow's resale prices in (186)–(189). Substituting (186)–(187) into (190)–(191) yields a linear system in (P_H, P_L). Under $M^F > 0$ this system has a unique solution, which can be written compactly as

$$\begin{pmatrix} P_H^* \\ P_L^* \end{pmatrix} = K^F \begin{pmatrix} D_H^e \\ D_L^e \end{pmatrix}$$

Step 3: Verify optimality (case-by-case). I verify that, at (P_H^*, P_L^*), the conjectured allocation is optimal given (82).

Case 1: $s_t = H$ and $s_{t-1} = H$. Type A enters with $S_{t-1}^A = 1$. Under the conjectured full-specialization allocation, a marginal share held in state H is valued by type A at Q_H^A. Since market clearing on the $L \to H$ transition pins down

$$Q_H^A = (1+\tau^F)P_H^*$$

and $(1+\tau^F)P_H^* > (1-\tau^F)P_H^*$, the sell condition in (82) fails. Hence type A does not sell and $S_t^A = 1$. Type B enters with $S_{t-1}^B = 0$. He does not buy provided

$$Q_H^B \leq (1+\tau^F)P_H^*$$

This inequality is implied by assumption (ii) of Proposition 5. Indeed, the sell incentive for type B on the $L \to H$ transition (Case 2) is

$$Q_H^B \leq (1-\tau^F)P_H^*$$

Because $(1-\tau^F)P_H^* < (1+\tau^F)P_H^*$, the latter condition is strictly stronger and therefore sufficient to rule out buying in the present case. Hence type B does not buy and $S_t^B = 0$.

Case 2: $s_t = H$ and $s_{t-1} = L$. Type A enters with $S_{t-1}^A = 0$ and buys because $Q_H^A = (1+\tau^F)P_H^*$ holds with equality by market clearing. Type B enters with $S_{t-1}^B = 1$ and sells if

$$Q_H^B \;\leq\; (1-\tau^F)P_H^* \tag{192}$$

Using (188) and evaluating prices at $\mathbf{P}^*$, Q_H^B can be rewritten as

$$Q_H^B = \delta\Big[D_L + (1+\tau^F)P_L^* + \pi_{H\to H}^B\Big((D_H - D_L) + (1-\tau^F)P_H^* - (1+\tau^F)P_L^*\Big)\Big]$$

(192) is equivalent to the upper bound

$$\pi_{H\to H}^B \;\leq\; \bar{\pi}_{H\to H}^B \;\equiv\; \frac{\frac{(1-\tau^F)P_H^*}{\delta} - \Big(D_L + (1+\tau^F)P_L^*\Big)}{(D_H - D_L) + (1-\tau^F)P_H^* - (1+\tau^F)P_L^*} \tag{193}$$

Under the maintained upper bound on τ^F, the denominator is positive. Hence, a sufficient condition for trade in the $L \to H$ transition is that $\pi_{H\to H}^B$ lies below the threshold $\bar{\pi}_{H\to H}^B$ in (193). If $\bar{\pi}_{H\to H}^B < 0$ then no $\pi_{H\to H}^B \in [0,1]$ can satisfy the sell condition, so equilibrium would fail. Since the denominator in (193) is positive, $\bar{\pi}_{H\to H}^B \geq 0$ is equivalent to

$$D_L \leq \Big(\frac{1-\tau^F}{\delta}\Big)P_H^* - (1+\tau^F)P_L^*$$

so that the numerator of (193) is positive. If $\bar{\pi}^B_{H\to H} > 1$, can safely be ignored.

Case 3: $s_t = L$ *and* $s_{t-1} = L$. Type B enters with $S^B_{t-1} = 1$. Market clearing on the $H \to L$ transition (Case 4) pins down

$$Q^B_L = (1+\tau^F)P^*_L$$

Since $(1+\tau^F)P^*_L > (1-\tau^F)P^*_L$, the sell condition in (82) fails for type B, so he does not sell and $S^B_t = 1$. Type A enters with $S^A_{t-1} = 0$. She does not buy provided

$$Q^A_L \le (1+\tau^F)P^*_L$$

This inequality is implied by the stronger condition that guarantees trade on the $H \to L$ transition (Case 4), namely

$$Q^A_L \le (1-\tau^F)P^*_L$$

Because $(1-\tau^F)P^*_L < (1+\tau^F)P^*_L$, the latter condition is sufficient to rule out buying here. Hence type A does not buy and $S^A_t = 0$.

Case 4: $s_t = L$ *and* $s_{t-1} = H$. Type B enters with $S^B_{t-1} = 0$ and buys because $Q^B_L = (1+\tau^F)P^*_L$ holds with equality by market clearing. Type A enters with $S^A_{t-1} = 1$ and sells if

$$Q^A_L \;\le\; (1-\tau^F)P^*_L \tag{194}$$

Using (189) and evaluating prices at $\mathbf{P}^*$, one can write

$$Q^A_L = \delta\Big[D_H + (1+\tau^F)P^*_H + \pi^A_{L\to L}\Big((D_L - D_H) + (1-\tau^F)P^*_L - (1+\tau^F)P^*_H\Big)\Big]$$

where $\pi^A_{L\to L} = 1 - \pi^A_{L\to H}$. (194) is equivalent to an upper bound on $\pi^A_{L\to H}$:

$$\pi^A_{L\to H} \;\le\; \bar{\pi}^A_{L\to H}(\tau^F) \;\equiv\; \frac{\frac{(1-\tau^F)(1-\delta)}{\delta}P^*_L - D_L}{(D_H - D_L) + (1+\tau^F)P^*_H - (1-\tau^F)P^*_L} \tag{195}$$

Under the maintained parameter restrictions used in the equilibrium construction the denominator is positive. Then, $\bar{\pi}^A_{L\to H}(\tau^F) \ge 0$ is equivalent to

$$\frac{(1-\tau^F)(1-\delta)}{\delta}\,P^*_L \;\ge\; D_L$$

If $\bar{\pi}^A_{L\to H}(\tau^F) > 1$, the limit can safely be ignored. Under these conditions, the bound in (195) lies in $[0,1]$, so there exists $\pi^A_{L\to H} \in [0,1]$ satisfying (194). Hence, a sufficient condition for trade in the $H \to L$ transition is that $\pi^A_{L\to H}$ lies below the threshold $\bar{\pi}^A_{L\to H}(\tau^F)$ in (195). Then, type A sells in (194) and the conjectured allocation in state L is optimal.

Thus $(S^A_t, S^B_t) = (0,1)$. In all histories, the conjectured full-specialization allocation clears the market and is individually optimal at the stationary prices (P^*_H, P^*_L). Therefore,

a competitive B-REE with (88) exists. ■

Proof of Proposition 6

Proof. Define the cross-state price gap in levels

$$\Delta^F(\tau^F) \equiv P_H^*(\tau^F) - P_L^*(\tau^F)$$

Since $P_t^* \in \{P_H^*, P_L^*\}$ with stationary probability μ on H and $(1-\mu)$ on L, the variance reads as

$$\mathrm{Var}(P_t^*) = \mu(1-\mu)\Delta^F(\tau^F)^2 \tag{196}$$

and hence

$$\frac{d}{d\tau^F}\mathrm{Var}(P_t^*) = 2\mu(1-\mu)\,\Delta^F(\tau^F)\,\frac{d\Delta^F}{d\tau^F} \tag{197}$$

Since $P_H^* > P_L^*$, we have $\Delta^F(\tau^F) > 0$, so the sign of $\frac{d}{d\tau^F}\mathrm{Var}(P_t^*)$ is pinned down by $\frac{d\Delta^F}{d\tau^F}$.

Equilibrium state prices can be written as

$$P_H^*(\tau^F) = \frac{\delta}{M^F(\tau^F)}\Big(d(\tau^F)D_H^e + b(\tau^F)D_L^e\Big)$$

$$P_L^*(\tau^F) = \frac{\delta}{M^F(\tau^F)}\Big(c(\tau^F)D_H^e + a(\tau^F)D_L^e\Big)$$

with

$$\begin{aligned} a(\tau^F) &= (1+\tau^F) - \delta\pi_A \\ d(\tau^F) &= (1+\tau^F) - \delta\pi_B \\ b(\tau^F) &= \delta(1-\pi_A)(1-\tau^F) \\ c(\tau^F) &= \delta(1-\pi_B)(1-\tau^F) \end{aligned}$$

and

$$M^F(\tau^F) = a(\tau^F)d(\tau^F) - b(\tau^F)c(\tau^F) = \big((1+\tau^F)-\delta\pi_A\big)\big((1+\tau^F)-\delta\pi_B\big) - \delta^2(1-\pi_A)(1-\pi_B)(1-\tau^F)^2 \tag{198}$$

Let

$$A(\tau^F) \equiv d(\tau^F)D_H^e + b(\tau^F)D_L^e, \qquad B(\tau^F) \equiv c(\tau^F)D_H^e + a(\tau^F)D_L^e.$$

Then

$$\Delta^F(\tau^F) = \frac{\delta}{M^F(\tau^F)}\Big(A(\tau^F) - B(\tau^F)\Big)$$

Differentiating gives (for given beliefs):

$$\frac{\partial \Delta^F}{\partial \tau^F} = \frac{\delta}{\big(M^F(\tau^F)\big)^2}\Big\{\big(A'(\tau^F) - B'(\tau^F)\big)M^F(\tau^F) - \big(A(\tau^F) - B(\tau^F)\big)\big(M^F\big)'(\tau^F)\Big\} \quad (199)$$

Step 1: Signs of A' and B'. Using $d' = a' = 1$, $b' = -\delta(1-\pi_A)$, and $c' = -\delta(1-\pi_B)$,

$$A'(\tau^F) = D_H^e - \delta(1-\pi_A)D_L^e, \qquad B'(\tau^F) = D_L^e - \delta(1-\pi_B)D_H^e.$$

On the B–REE set characterized in Proposition 5, we have $M^F(\tau^F) > 0$ and $P_H^*(\tau^F) > P_L^*(\tau^F)$ (by condition (i) there together with $D_H > D_L$), hence

$$\Delta^F(\tau^F) = \frac{\delta}{M^F(\tau^F)}\big(A(\tau^F) - B(\tau^F)\big) > 0, \qquad \Rightarrow \qquad A(\tau^F) - B(\tau^F) > 0.$$

Moreover, under the maintained belief ordering $\pi_A \geq \pi_B$ and $D_H > D_L$, expected dividends are ordered:

$$D_H^e \;=\; D_L + \pi_A(D_H - D_L) \;\geq\; D_L + (\pi_B)(D_H - D_L) \;=\; D_L^e.$$

Therefore,

$$A'(\tau^F) - B'(\tau^F) = \big(1+\delta(1-\pi_B)\big)D_H^e - \big(1+\delta(1-\pi_A)\big)D_L^e \;\geq\; \big(1+\delta(1-\pi_A)\big)\big(D_H^e - D_L^e\big) \;\geq\; 0,$$

so $A'(\tau^F) - B'(\tau^F) \geq 0$.

Step 2: Sign of $(M^F)'$. From (198), using $a' = d' = 1$ and $(bc)' = b'c + bc'$,

$$\big(M^F\big)'(\tau^F) = a(\tau^F) + d(\tau^F) + 2\delta^2(1-\pi_A)(1-\pi_B)(1-\tau^F)$$

On the maintained parameter set $\delta \in (0,1)$, $\pi_A, \pi_B \in [0,1]$, and $\tau^F \in [0,1]$, the last term is nonnegative. Moreover,

$$a(\tau^F) = (1+\tau^F) - \delta\pi_A \;\geq\; 1-\delta \;>\; 0, \qquad d(\tau^F) = (1+\tau^F) - \delta\pi_B \;\geq\; 1-\delta \;>\; 0,$$

hence $\big(M^F\big)'(\tau^F) > 0$.

Step 3: Sufficient condition for $\frac{\partial \Delta^F}{\partial \tau^F} \geq 0$. Equation (199) implies that $\frac{\partial \Delta^F}{\partial \tau^F} \geq 0$ if and only if

$$\big(A'(\tau^F) - B'(\tau^F)\big)M^F(\tau^F) \;\geq\; \big(A(\tau^F) - B(\tau^F)\big)\big(M^F\big)'(\tau^F)$$

or, equivalently,

$$\frac{A'(\tau^F) - B'(\tau^F)}{A(\tau^F) - B(\tau^F)} \;\geq\; \frac{\big(M^F\big)'(\tau^F)}{M^F(\tau^F)} \tag{200}$$

On the B–REE set characterized in Proposition 5, we have $M^F(\tau^F) > 0$ and $A(\tau^F) - B(\tau^F) > 0$ (since $P_H^*(\tau^F) > P_L^*(\tau^F)$), so both sides of (200) are well-defined. Moreover, Step 2 implies $\big(M^F\big)'(\tau^F) > 0$.

A convenient one-sided sufficient condition can be expressed directly as a lower bound on π_A. Let

$$C_0(\pi_B) \equiv (1-\delta)\Big(\frac{1}{\delta(1-\pi_B)} - 1\Big)$$

$$C_1(\tau^F, \pi_B) \equiv \frac{2\big((1+\tau^F) - \delta\pi_B\big) + 2\delta^2(1-\pi_B)^2(1-\tau^F)}{\big((1+\tau^F) - \delta\big)^2 - \delta^2(1-\pi_B)^2(1-\tau^F)^2}$$

By Proposition 5, the feasibility restriction $\tau^F \leq \bar{\tau}^F(\mathbf{P}^*(\tau^F))$ ensures that the relevant denominators are positive when evaluated at the equilibrium price vector. I further restrict attention to the subset of this region where $1 - \tau^F C_1(\tau^F, \pi_B) > 0$ so that the rearrangement of (200) preserves inequality direction. Under this sign condition, (200) is implied by the one-sided bound

$$\pi_A \;\geq\; \underline{\pi}_A^F(\tau^F, \pi_B) \equiv 1 - \frac{1}{\delta}\left[\frac{1}{\delta(1-\pi_B)} - \frac{C_1(\tau^F, \pi_B)\, C_0(\pi_B)}{1 - \tau^F C_1(\tau^F, \pi_B)}\right] \tag{201}$$

Condition (201) ensures that the sensitivity of the wedge $\Delta^F(\tau^F)$ to τ^F dominates the sensitivity of the common scaling term $M^F(\tau^F)$.

Boundary effect. When a trading constraint binds, a higher transaction tax reduces the seller's net-of-tax proceeds in the relevant state. To keep trade feasible on the binding frontier, beliefs must adjust so that the state price in that same state increases enough to compensate the seller for the tax wedge. Along frontier (ii) $\pi^B_{H\to H} = \pi^A_{H\to H} - \kappa^F_H(\tau^F)$, this compensation operates through an increase in the high-state valuation P_H^*, which tends to widen the price gap $\Delta^F \equiv P_H^* - P_L^*$. Along frontier (iii) $\pi^A_{L\to H} = \pi^B_{L\to H} - \kappa^F_L(\tau^F)$, the analogous compensation operates through an increase in the low-state valuation P_L^*, which tends to compress Δ^F. Because the binding frontier depends on which constraint is active and because the two frontiers move different components of $\mathbf{P}^*$, the net frontier contribution to $d\Delta^F/d\tau^F$ is in general ambiguous. For this reason, I do not seek a general sign result for the frontier component. Instead, I assume that the direct effect of the tax on the gap, together with the positive belief adjustment on frontier (ii) when it binds, dominates the compensating mechanism on frontier (iii) that raises P_L^* by increasing recovery optimism and therefore compresses Δ^F, preserving the (positive) sign of the total derivative. ■

A.8 Proof of Lemma 4

Proof. Let the loading matrix be

$$\Lambda^F(\tau^F) \;\equiv\; \frac{\delta}{1+\tau^F}\,\mathcal{B}^F(\tau^F), \qquad \Lambda^F_s(\tau^F) \text{ the } s\text{-row}.$$

For each state s,

$$\mathrm{Var}(P_{st}) = \;\Lambda^F_s\,\Sigma_\beta\,(\Lambda^F_s)' \tag{202}$$

Differentiating it with respect to τ^F,

$$\frac{d}{d\tau^F}\mathrm{Var}(P_{st}) = 2\left(\frac{d\Lambda^F_s}{d\tau^F}\right)\Sigma_\beta\,(\Lambda^F_s)' \;+\; \Lambda^F_s\,\frac{d\Sigma_\beta}{d\tau^F}\,(\Lambda^F_s)'. \tag{203}$$

Effect of τ^F on loadings

Hold $\Sigma_\beta = \begin{pmatrix}\sigma_{HH} & \sigma_{HL}\\ \sigma_{HL} & \sigma_{LL}\end{pmatrix}$ fixed. From (94), each element of $\Lambda^F(\tau^F)$ is strictly decreasing on $\tau^F \in (0,1)$ (the diagonal entries scale as $(1+\tau^F)^{-1}$, while off-diagonals scale as $(1+\tau^F)^{-1}(1-\tau^F)$). Let $\Lambda^F_s = (\lambda^F_{s1}, \lambda^F_{s2})$. Then

$$\frac{1}{2}\frac{d}{d\tau^F}\mathrm{Var}(P_{st})\Big|_{\Sigma_\beta} = \sigma_{HH}\lambda^F_{s1}\frac{d\lambda^F_{s1}}{d\tau^F} + \sigma_{LL}\lambda^F_{s2}\frac{d\lambda^F_{s2}}{d\tau^F} + \sigma_{HL}\left(\lambda^F_{s2}\frac{d\lambda^F_{s1}}{d\tau^F} + \lambda^F_{s1}\frac{d\lambda^F_{s2}}{d\tau^F}\right). \tag{204}$$

For $\pi_A, \pi_B \in (0,1)$ and $\tau^F \in (0,1)$, we have $\lambda^F_{s1}, \lambda^F_{s2} > 0$ and $d\lambda^F_{s1}/d\tau^F < 0$, $d\lambda^F_{s2}/d\tau^F < 0$. Hence, if $\sigma_{HL} \geq 0$, each term in (204) is negative and

$$\frac{d}{d\tau^F}\mathrm{Var}(P_{st})\Big|_{\Sigma_\beta} < 0 \qquad \text{for each } s \in \{H, L\}. \tag{205}$$

Effect of τ^F on the beliefs covariance matrix

Adopt the same mean–coefficient approximation as in the CGT proof: replace $\mathrm{diag}(\mathbf{I}_{t-1})$ by its stationary expectation

$$\Lambda \equiv \mathbb{E}[\mathrm{diag}(\mathbf{I}_{t-1})] = \mathrm{diag}(\lambda_H, \lambda_L).$$

Using the FTT pricing rule (93) and the loading matrix $\Lambda^F(\tau^F) \equiv \frac{\delta}{1+\tau^F}\mathcal{B}^F(\tau^F)$, the mean–coefficient belief recursion is the VAR(1)

$$\boldsymbol{\beta}_t = \boldsymbol{d}^F + \boldsymbol{U}^F(\tau^F)\,\boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t, \qquad \mathbb{E}[\boldsymbol{\eta}_t] = \mathbf{0}, \qquad \mathrm{Var}(\boldsymbol{\eta}_t) = \Sigma_u = \sigma_b^2\, G\Lambda G, \tag{206}$$

with $G = \mathrm{diag}(g_H, g_L)$ and

$$\boldsymbol{d}^F = G\Lambda\,\delta\,\boldsymbol{D}^e, \qquad \boldsymbol{U}^F(\tau^F) = I_2 - G\Lambda + G\Lambda\,\Lambda^F(\tau^F). \tag{207}$$

Since $\boldsymbol{U}^F(\tau^F)$ is mean–stable ($\rho(\boldsymbol{U}^F(\tau^F)) < 1$)[57], Σ_β solves

$$\Sigma_\beta = \boldsymbol{U}^F \, \Sigma_\beta \, (\boldsymbol{U}^F)^\top + \Sigma_u. \tag{208}$$

As in the CGT proof, since Σ_u is diagonal with positive entries and $\boldsymbol{U}^F$ has nonnegative entries, the Neumann representation implies Σ_β has nonnegative entries; I omit the repetition and refer to the corresponding step in the CGT proof.

Differentiating (208) w.r.t. τ^F (noting Σ_u does not depend on τ^F) yields the Lyapunov equation

$$\Sigma_\tau^F = \boldsymbol{U}^F \, \Sigma_\tau^F \, (\boldsymbol{U}^F)^\top + \boldsymbol{W}^F, \qquad \Sigma_\tau^F \equiv \frac{d\Sigma_\beta}{d\tau^F}, \qquad \boldsymbol{W}^F \equiv \boldsymbol{U}_\tau^F \Sigma_\beta (\boldsymbol{U}^F)^\top + \boldsymbol{U}^F \Sigma_\beta (\boldsymbol{U}_\tau^F)^\top, \tag{209}$$

where

$$\boldsymbol{U}_\tau^F \equiv \frac{d\boldsymbol{U}^F}{d\tau^F} = G\Lambda \, \frac{d\Lambda^F(\tau^F)}{d\tau^F}. \tag{210}$$

Mean stability implies the Neumann-series representation

$$\Sigma_\tau^F = \sum_{k=0}^{\infty} (\boldsymbol{U}^F)^k \, \boldsymbol{W}^F \, \big((\boldsymbol{U}^F)^\top\big)^k. \tag{211}$$

Sign of the Lyapunov innovation $\boldsymbol{W}^F$. By construction, $\boldsymbol{U}^F$ has nonnegative entries. Moreover, by the monotonicity of the FTT loadings established above, $\Lambda^F(\tau^F)$ is elementwise decreasing in $\tau^F \in (0,1)$, hence $d\Lambda^F(\tau^F)/d\tau^F \leq \mathbf{0}$ elementwise. Since $G\Lambda$ has positive diagonal entries, (210) implies

$$\boldsymbol{U}_\tau^F \leq \mathbf{0} \quad \text{elementwise (with strict inequalities generically).} \tag{212}$$

Because Σ_β has nonnegative entries and $(\boldsymbol{U}^F)^\top$ has nonnegative entries, the two terms defining $\boldsymbol{W}^F$ are elementwise nonpositive:

$$\boldsymbol{U}_\tau^F \Sigma_\beta (\boldsymbol{U}^F)^\top \leq \mathbf{0}, \qquad \boldsymbol{U}^F \Sigma_\beta (\boldsymbol{U}_\tau^F)^\top \leq \mathbf{0},$$

hence

$$\boldsymbol{W}^F \leq \mathbf{0} \quad \text{elementwise.} \tag{213}$$

Furthermore, if $\boldsymbol{U}_\tau^F$ has at least one strictly negative entry (generic under (212)) and Σ_β is not the zero matrix, then $\boldsymbol{W}^F$ has at least one strictly negative entry.

Quadratic forms for both states. Fix $s \in \{H, L\}$ and write the s-row of $\Lambda^F(\tau^F)$ as $\Lambda_s^F(\tau^F) = (\lambda_{s1}^F, \lambda_{s2}^F)$. For $\pi_A, \pi_B \in (0,1)$ and $\tau^F \in (0,1)$, we have $\lambda_{s1}^F, \lambda_{s2}^F > 0$, so Λ_s^F has

[57] Mean–stability follows because $\boldsymbol{U}^F(\tau^F)$ has nonnegative entries and its row sums satisfy $\sum_j U_{sj}^F = 1 - g_s\lambda_s\big(1 - \sum_j \Lambda_{sj}^F(\tau^F)\big) < 1$ for $\tau^F > 0$ (and $g_s\lambda_s > 0$), hence $\rho(\boldsymbol{U}^F(\tau^F)) \leq \|\boldsymbol{U}^F(\tau^F)\|_\infty < 1$.

strictly positive entries. Combining (211) with $\boldsymbol{U}^F \geq \mathbf{0}$ and (213) implies

$$\Sigma_\tau^F \leq \mathbf{0} \quad \text{elementwise.}$$

Since each term in the Neumann sum is elementwise nonpositive, it follows that the scalar quadratic form is nonpositive:

$$\Lambda_s^F \, \Sigma_\tau^F \, (\Lambda_s^F)' \leq 0.$$

Finally, strict negativity holds because $\boldsymbol{W}^F \not\equiv 0$ implies that the $k = 0$ term in (211) already contributes strictly:

$$\Lambda_s^F \, \Sigma_\tau^F \, (\Lambda_s^F)' = \sum_{k=0}^{\infty} \Lambda_s^F (\boldsymbol{U}^F)^k \, \boldsymbol{W}^F \left((\boldsymbol{U}^F)^\top\right)^k (\Lambda_s^F)' \; < \; 0,$$

because $\Lambda_s^F \boldsymbol{W}^F (\Lambda_s^F)' < 0$ whenever $\boldsymbol{W}^F$ has at least one strictly negative entry. Therefore, for both $s \in \{H, L\}$,

$$\Lambda_s^F \, \frac{d\Sigma_\beta}{d\tau^F} \, (\Lambda_s^F)' < 0 \tag{214}$$

■

E. Details of the Quantitative Model

E.1. Derivation of the marginal continuation value

Denote $\psi(z_t^i) \equiv \frac{1}{P_t} \delta \mathbb{E}^{\mathcal{P}^i} \left[\frac{\partial V^i(z_{t+1}^i)}{\partial S_t^i} \,\middle|\, z_t^i \right] = \frac{Q_t^i}{P_t}$. Thus, $\psi(z_t^i)$ summarizes investor i's discounted subjective marginal value of carrying one additional share into $t+1$ relative to the current price. Using the generalized envelope theorem of Milgrom and Segal (2002) together with the chain rule, the marginal continuation value can be written as

$$\begin{aligned} \psi(z_t^i) = \mathbb{E}_t^{\mathcal{P}^i} \Bigg[& \delta (C_{t+1}^i)^{-\gamma} \frac{1}{P_t} \Big(D_{t+1}(1 - \tau_{t+1}^D) + P_{t+1} - \tau_{t+1}^{CG} \big(\max\{0,\, P_{t+1} - P_{b,t+1}^i\} \, \mathbb{1}\{S_{t+1}^i < \hat{S}_{t+1}^i\} \big) \Big) \\ & + \frac{\delta}{P_t} (F_{1,t+1}^i + F_{2,t+1}^i) \Bigg] \end{aligned} \tag{215}$$

The terms $F_{1,t+1}^i, F_{2,t+1}^i$ capture, respectively, the marginal impact of holdings on carryover losses and basis price.

The direct marginal tax wedge from an extra share is proportional to the per-share gain $\max\{0, P_{t+1} - P_{b,t+1}^i\}$ and applies only when the marginal share is sold in the taxed region, i.e. when $S_{t+1}^i < \hat{S}_{t+1}^i$. $\hat{S}_{t+1}^i$ denotes the smallest post-trade position at $t+1$ that

exhausts the loss stock exactly. Under the loss-offset rule, this benchmark is

$$\hat{S}^i_{t+1} = \begin{cases} S^i_t \max\left\{0,\, 1 - \dfrac{L^i_{t+1}}{P_{t+1} - P^i_{b,t+1}}\right\} & \text{if} \quad P_{t+1} > P^i_{b,t+1} \\ 0 & \text{otherwise} \end{cases} \tag{216}$$

The logic is that sales are tax-free until losses are exhausted; beyond that (i.e., for $S^i_{t+1} < \hat{S}^i_{t+1}$) the usual capital-gains wedge applies.

Moreover, an extra share at t affects the tax base at $t+1$ through the loss stock L^i_{t+1}. Because the loss L^i_{t+1} is a state variable, its marginal value satisfies a recursive envelope condition. Let $V^i_L(z^i_{t+1}) \equiv \frac{\partial V^i(z^i_{t+1})}{\partial L^i_{t+1}}$. Differentiating the Bellman equation at $t+1$ yields

$$V^i_L(z^i_{t+1}) = (C^i_{t+1})^{-\gamma} \frac{\partial C^i_{t+1}}{\partial L^i_{t+1}} + \delta\, \mathbb{E}^{\mathcal{P}^i}_{t+1}\left[V^i_L(z^i_{t+2}) \frac{\partial L^i_{t+2}}{\partial L^i_{t+1}}\right] \tag{217}$$

where the first term captures the direct effect of a marginal increase in L^i_{t+1} on the period-$t+1$ budget through the tax schedule, and the second term captures the continuation effect operating through the law of motion of the loss state. A marginal increase in the loss stock raises $t+1$ consumption one-for-one (at rate τ^{CG}_{t+1}) whenever the household sells beyond the loss-exhaustion threshold:

$$\frac{\partial C^i_{t+1}}{\partial L^i_{t+1}} = \tau^{CG}_{t+1} S^i_t \mathbb{1}\{S^i_{t+1} < \hat{S}^i_{t+1}\}$$

Let the marginal continuation effect of S^i_t through the loss state be

$$\frac{\partial V^i(z^i_{t+1})}{\partial L^i_{t+1}} \frac{\partial L^i_{t+1}}{\partial S^i_t} \equiv F^i_{1,t+1} \tag{218}$$

Under the loss carryover law of motion (40), the derivative is piecewise:

$$\frac{\partial L^i_{t+1}}{\partial S^i_t} = \mathbb{1}\{L^i_t S^i_{t-1} - (P_t - P^i_{b,t}) \max\{0, S^i_{t-1} - S^i_t\} > 0\} \begin{cases} -\dfrac{L^i_t S^i_{t-1}}{(S^i_t)^2} & S^i_t \geq S^i_{t-1}, \\ \dfrac{\left((P_t - P^i_{b,t}) - L^i_t\right) S^i_{t-1}}{(S^i_t)^2} & S^i_t < S^i_{t-1} \end{cases} \tag{219}$$

This term captures how an extra holding at t changes next period's per-share loss stock through dilution (the $1/S^i_t$ term) and, when selling at t, through the numerator in the indicator function.

A similar logic applies to the basis price: an extra share at t also affects the tax base

at $t+1$ through the average-cost basis $P^i_{b,t+1}$:

$$\frac{\partial V^i(z^i_{t+1})}{\partial P^i_{b,t+1}} \frac{\partial P^i_{b,t+1}}{\partial S^i_t} \equiv F^i_{2,t+1} \tag{220}$$

where

$$V^i_B(z^i_{t+1}) = (C^i_{t+1})^{-\gamma} \frac{\partial C^i_{t+1}}{\partial P^i_{b,t+1}} + \delta \, \mathbb{E}^{\mathcal{P}^i}_{t+1} \left[V^i_B(z^i_{t+2}) \frac{\partial P^i_{b,t+2}}{\partial P^i_{b,t+1}} \right] \tag{221}$$

with

$$\frac{\partial C^i_{t+1}}{\partial P^i_{b,t+1}} \;=\; \tau^{CG}_{t+1} \max\{0,\, S^i_t - S^i_{t+1}\} \, \mathbb{1}\{S^i_{t+1} < \hat{S}^i_{t+1}\}$$

that is, increasing the average basis reduces realized gains per share, hence raises consumption, but only when the household is in the taxed region (net taxable base strictly positive). In turn, under the average-cost basis rule (41), for $S^i_t > 0$,

$$\frac{\partial P^i_{b,t+1}}{\partial S^i_t} = \mathbb{1}\{S^i_t > S^i_{t-1}\} \, \frac{(P_t - P^i_{b,t}) \, S^i_{t-1}}{(S^i_t)^2} \tag{222}$$

This term is active only when the investor buys at t (so the average basis is updated) and it captures how an additional share affects future taxable gains by shifting the average-cost basis.

In the quantitative implementation, I adopt the myopic-envelope approximation $V_L(z_t) \approx u'(C_t)\, \partial C_t / \partial L_t$ and $V_B(z_t) \approx u'(C_t)\, \partial C_t / \partial P_{b,t}$, which retains the leading-order effect of tax assets on the marginal valuation of the marginal share. The omitted continuation components matter only when the realization kink is active and are further scaled by $\partial L_{t+1}/\partial S_t$ and $\partial P_{b,t+1}/\partial S_t$, which inherit the S_t^{-2} dilution terms implied by average-cost accounting. Second, and more directly, even setting the full leading-order contributions F_1 and F_2 to zero leaves simulated policy functions virtually unchanged, confirming that the indirect tax-state channels are quantitatively second-order for equilibrium allocations in the baseline parameterization.

E.2. Dealing with inequality constraints

The optimal allocations must satisfy the sequence of budget constraints and this list of conditions:

$$
\begin{aligned}
C_t^i \geq \psi_t^{-1/\gamma} \quad & \text{and} \quad S_t^i = \bar{S}; \\
C_t^i = \psi_t^{-1/\gamma} \quad & \text{and} \quad S_t^i \in [S_{t-1}^i, \bar{S}] \cup [\hat{S}_t^i, S_{t-1}^i]; \\
\psi_t^{-1/\gamma} \geq C_t^i \geq \left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \quad & \text{and} \quad S_t^i = \hat{S}_t^i \ \text{ if } \ P_t > P_{bt}^i; \\
C_t^i = \left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \quad & \text{and} \quad S_t^i \in [0, \hat{S}_t^i] \\
\left(\frac{\psi_t}{\kappa_t^i}\right)^{-1/\gamma} \geq C_t^i \quad & \text{and} \quad S_t^i = 0;
\end{aligned}
\tag{223}
$$

where $\psi_t = \psi(z_t^i)$ and $\kappa_t = 1 - \tau_t^{CG} \max\{\left(1 - \frac{P_{bt}^i}{P_t}\right), 0\}$.

The following algorithm, in the spirit of Marcet and Singleton (1999), ensures that all the optimality conditions summarized by the KKT conditions are satisfied:

1. Start in the non-selling region. Set $C_0^i = \psi_t^{-1/\gamma}$ and using the budget constraint, $S_0^i = \frac{X_t^i - C_0^i}{P_t}$, where $X_t^i = W_t + ((1 - \tau_t^D)D_t + P_t)S_{t-1}^i + T_t^i$.

2. If no tax at the margin, i.e., either buy/hold $S_0^i \geq S_{t-1}^i$, no gains $(P_t - P_{bt}^i) \leq 0$ or candidate sale is within losses $S_0^i \geq \hat{S}_t^i$; then, accept the candidate $(C_t^i, S_t^i) = (C_0^i, S_0^i)$ and go to step 5.

3. Otherwise, taxed sell candidate: $S_0^i < \hat{S}_t^i$, and $P_t > P_{bt}^i$. Set the sell-side equality consumption $C_1^i = \left(\frac{\psi_t}{\kappa_t}\right)^{-1/\gamma}$ and compute S_1^i from the budget constraint: $S_1^i = \frac{X_t^i - \tau_t^{CG}(P_t - P_{bt}^i - L_t^i)S_{t-1}^i - C_1^i}{P_t - \tau_t^{CG}(P_t - P_{bt}^i)}$.

4. Selling feasibility check.

   (a) If $S_1^i \leq \hat{S}_t^i$, set $(C_t^i, S_t^i) = (C_1^i, S_1^i)$ and go to step 5.

   (b) If $S_1^i > \hat{S}_t^i$, the sell-equality lies to the right of the kink (implies selling less than needed to reach the taxed region). The optimum is the kink: $(C_t^i, S_t^i) = (X_t^i - P_t\hat{S}_t^i, \hat{S}_t^i)$.

5. Bounds projection.

   (a) If $S_t^i > \bar{S}$, set $S_t^i = \bar{S}$, and using the budget constraint $C_t^i = X_t^i - P_t\bar{S}$.

   (b) If $S_t^i < 0$, set $S_t^i = 0$, and using the budget constraint, $C_t^i = X_t^i - \tau_t^{CG} \max\Big\{ \max\{0, S_{t-1}^i - S_t^i\}(P_t - P_{bt}^i) - L_t^i S_{t-1}^i, 0 \Big\}$.

Accrual-based taxation

Let $\Delta P_{t+1} \equiv P_{t+1} - P_t$ and recall the accrual-loss transition

$$L^i_{t+1} = \frac{S^i_{t-1}}{S^i_t} \max\{L^i_t - (P_t - P_{t-1}), 0\}$$

At $t+1$, total accrual-tax payments are

$$\tau^A_{t+1}\, S^i_t\, \max\{0, \Delta P_{t+1} - L^i_{t+1}\}$$

Hence the marginal (one-share) payoff from increasing S^i_t has two components: (i) the direct effect on next-period resources, and (ii) the indirect effect through the induced change in L^i_{t+1}. Writing the Bellman equation as $V^i_t(z^i_t) = \max_{C^i_t, S^i_t}\{u(C^i_t) + \delta\, \mathbb{E}^{\mathcal{P}^i}_t[V^i_{t+1}(z^i_{t+1})]\}$, the interior FOC for S^i_t is

$$u'(C^i_t)\, P_t = \delta\, \mathbb{E}^{\mathcal{P}^i}_t \left[V^i_{W,t+1} \frac{\partial W^i_{t+1}}{\partial S^i_t} \;+\; V^i_{L,t+1} \frac{\partial L^i_{t+1}}{\partial S^i_t} \right], \tag{224}$$

where $V^i_{W,t+1} \equiv \partial V^i_{t+1}/\partial W^i_{t+1}$ and $V^i_{L,t+1} \equiv \partial V^i_{t+1}/\partial L^i_{t+1}$.

From the losses law of motion,

$$\frac{\partial L^i_{t+1}}{\partial S^i_t} = -\frac{S^i_{t-1}}{(S^i_t)^2} \max\{L^i_t - (P_t - P_{t-1}), 0\} = -\frac{L^i_{t+1}}{S^i_t} \tag{225}$$

Moreover, differentiating the tax term,

$$\frac{\partial}{\partial S^i_t}\Big[S^i_t \max\{0, \Delta P_{t+1} - L^i_{t+1}\} \Big] = \Delta P_{t+1}\, \mathbf{1}\{\Delta P_{t+1} > L^i_{t+1}\}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function.[58] Therefore the direct marginal effect of S^i_t on next-period resources is

$$\frac{\partial W^i_{t+1}}{\partial S^i_t} = (1 - \tau^D_{t+1}) D_{t+1} + P_{t+1} - \tau^A_{t+1}\, \Delta P_{t+1}\, \mathbf{1}\{\Delta P_{t+1} > L^i_{t+1}\} \tag{226}$$

Using the envelope condition $V^i_{W,t+1} = u'(C^i_{t+1})$ and substituting (225)–(226) into (224)

[58] This identity uses $\frac{\partial L^i_{t+1}}{\partial S^i_t} = -L^i_{t+1}/S^i_t$ and the product rule: $\partial[S \max\{0, \Delta P - L(S)\}]/\partial S = \max\{0, \Delta P - L\} + S\, \mathbf{1}\{\Delta P > L\}(-\partial L/\partial S)$. When $\Delta P > L$, this becomes $(\Delta P - L) + L = \Delta P$; otherwise it is 0.

yields

$$(C_t^i)^{-\gamma} = \mathbb{E}_t^{\mathcal{P}^i}\Bigg[\delta\,(C_{t+1}^i)^{-\gamma}\,\frac{1}{P_t}\Big((1-\tau_{t+1}^D)D_{t+1}+P_{t+1}-\tau_{t+1}^A\,\Delta P_{t+1}\,\mathbf{1}\{\Delta P_{t+1}>L_{t+1}^i\}\Big) - \delta\,\frac{1}{P_t}\,V_{L,t+1}^i\,\frac{L_{t+1}^i}{S_t^i}\Bigg] \equiv \psi^A(z_{A,t}^i) \tag{227}$$

Since L_t^i affects next period only through L_{t+1}^i, one has

$$V_{L,t}^i = u'\big(C_t^i\big)\;\tau_t^A\,S_{t-1}^i\,\mathbf{1}\big\{P_t-P_{t-1}>L_t^i\big\}\;+\;\delta\,\mathbb{E}_t^{\mathcal{P}^i}\left[V_{L,t+1}^i\,\frac{\partial L_{t+1}^i}{\partial L_t^i}\right]$$

with

$$\frac{\partial L_{t+1}^i}{\partial L_t^i} = \frac{S_{t-1}^i}{S_t^i}\,\mathbf{1}\big\{L_t^i>P_t-P_{t-1}\big\}$$

where the equality follows from (A.8). As in the case of realization-based taxation, in simulations I adopt the myopic-envelope approximation to $V_{L,t}^i$.

The algorithm to deal with the holding limits is a simplified version of the realization case. The KKT conditions boil down to

$$\begin{aligned} &C_t^i \geq (\psi_t^A)^{-1/\gamma} \text{and}\;\; S_t^i=\bar{S}\\ &C_t^i = (\psi_t^A)^{-1/\gamma} \text{and}\;\; S_t^i\in(0,\bar{S})\\ &C_t^i \leq (\psi_t^A)^{-1/\gamma} \text{and}\;\; S_t^i=0 \end{aligned} \tag{228}$$

1. Start in the non-binding region. Set $C_0^i=(\psi_t^A)^{-1/\gamma}$ and using the budget constraint, $S_0^i=\frac{X_t^i-C_0^i}{P_t}$, where $X_t^i=W_t+((1-\tau_t^D)D_t+P_t)S_{t-1}^i-\tau_t^A\max\{P_t-P_{t-1}-L_t^i,0\}S_{t-1}^i+T_t^i$.
2. If $S_0^i\in[0,1]$, accept the candidate $(C_t^i,S_t^i)=(C_0^i,S_0^i)$ and stop.
3. Bounds projection.
   (a) If $S_t^i>1$, set $S_t^i=1$, and using the budget constraint $C_t^i=X_t^i-P_t$.
   (b) If $S_t^i<0$, set $S_t^i=0$, and using the budget constraint, $C_t^i=X_t^i$.

Adding a Financial Transactions Tax

Consider the extended model discussed in Section 3 and include a trading cost. The budget constraint becomes:

$$\begin{aligned} C_t^i+P_tS_t^i =& W_t+((1-\tau_t^D)D_t+P_t)S_{t-1}^i-\tau_t^F|S_{t-1}^i-S_t^i|P_t\\ &-\tau_t^{CG}\max\Big\{\max\{0,S_{t-1}^i-S_t^i\}(P_t-P_{bt}^i)-L_t^iS_{t-1}^i,0\Big\}+T_t^i \end{aligned} \tag{229}$$

It follows that the discounted expected marginal value of carrying one additional share into $t+1$ relative to the current price has an extra term:

$$\psi^F(z_t^i) = \psi(z_t^i) - \mathbb{E}_t^{\mathcal{P}^i}\left[\delta(C_{t+1}^i)^{-\gamma}\frac{P_{t+1}}{P_t}\Big(\tau_{t+1}^F \mathbb{1}\{S_{t+1}^i < S_t^i\} - \tau_{t+1}^F \mathbb{1}\{S_{t+1}^i > S_t^i\}\Big)\right] \quad (230)$$

With an FTT, the investor faces an additional proportional wedge on traded value. Let $\psi_t^F \equiv \psi(z_t^i)^F$ as in the baseline model and define

$$\overline{\kappa}_t^F \equiv 1 + \tau_t^F$$

$$\underline{\kappa}_t^F \equiv 1 - \tau_t^F$$

$$\kappa_t^{F,i} \equiv (1 - \tau_t^F) \; - \; \tau_t^{CG} \max\Big\{1 - \frac{P_{bt}^i}{P_t},\, 0\Big\}$$

Thus, a marginal purchase costs $\overline{\kappa}_t^F P_t$, a marginal CGT tax-free sale yields $\underline{\kappa}_t^F P_t$, and a marginal taxable sale yields $\kappa_t^{F,i} P_t$. Because an FTT implies different marginal price/proceeds factors for purchases and sales, and CGT further splits the selling side into taxed versus tax-free marginal regimes, I state the KKT conditions separately by trading side:

Buying (weakly increasing shares).

$$C_t^i \geq \left(\frac{\psi_t^F}{\overline{\kappa}_t^F}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i = \bar{S}$$

$$C_t^i = \left(\frac{\psi_t^F}{\overline{\kappa}_t^F}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i \in (S_{t-1}^i, \bar{S})$$

Selling with a CGT tax-free marginal sale. If $P_t \leq P_{bt}^i$ (no embedded gains), then any marginal sale is CGT tax-free and the relevant sell-side wedge is $\underline{\kappa}_t^F$. If instead $P_t > P_{bt}^i$, the marginal sale is CGT tax-free when $S_t^i \in [\hat{S}_t^i, S_{t-1}^i]$. In either case,

$$C_t^i = \left(\frac{\psi_t^F}{\underline{\kappa}_t^F}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i \in (\hat{S}_t^i, S_{t-1}^i]$$

$$C_t^i \leq \left(\frac{\psi_t^F}{\underline{\kappa}_t^F}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i = \hat{S}_t^i$$

Selling below the loss-offset threshold (taxable marginal sale). If the marginal sale is taxable (equivalently $S_t^i \in [0, \hat{S}_t^i)$ when $P_t > P_{bt}^i$),

$$C_t^i = \left(\frac{\psi_t^F}{\kappa_t^{F,i}}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i \in (0, \hat{S}_t^i)$$

$$C_t^i \leq \left(\frac{\psi_t^F}{\kappa_t^{F,i}}\right)^{-1/\gamma} \quad \text{and} \quad S_t^i = 0$$

At $S_t^i = S_{t-1}^i$, no-trade obtains when consumption lies between the buy-side marginal

condition and the relevant sell-side marginal condition. Specifically:

$$\left(\frac{\psi_t^F}{\overline{\kappa}_t^F}\right)^{-1/\gamma} \geq C_t^i \geq \left(\frac{\psi_t^F}{\underline{\kappa}_t^F}\right)^{-1/\gamma} \quad \text{if a marginal sale would be CGT tax-free}$$

which holds whenever $P_t \leq P_{bt}^i$ or, if $P_t > P_{bt}^i$, a marginal decrease from S_{t-1}^i remains weakly above $\hat{S}_t^i$. Otherwise (i.e., $P_t > P_{bt}^i$ and a marginal decrease would fall below $\hat{S}_t^i$),

$$\left(\frac{\psi_t^F}{\overline{\kappa}_t^F}\right)^{-1/\gamma} \geq C_t^i \geq \left(\frac{\psi_t^F}{\kappa_t^{F,i}}\right)^{-1/\gamma} \quad \text{if a marginal sale would be taxable}$$

In words, the FTT replaces the purchase price P_t by the effective buy price $(1+\tau_t^F)P_t$, and it scales down marginal sale proceeds by the factor $(1-\tau_t^F)$. When the marginal sale is subject to capital-gains taxation (i.e., when $P_t > P_{bt}^i$ and the marginal reduction in S_t^i falls below the loss-offset threshold $\hat{S}_t^i$), proceeds are further reduced by the realization wedge $\tau_t^{CG}\max\{1 - P_{bt}^i/P_t, 0\}$. As a result, the no-trade region is characterized by ψ_t^F lying between the buy-side marginal condition and the relevant sell-side marginal condition: the tax-free wedge $(1-\tau_t^F)$ if the marginal sale is CGT-free, and the taxable wedge $\kappa_t^{F,i}$ otherwise. The portfolio bounds $S_t^i \in [0, \bar{S}]$ and the threshold $\hat{S}_t^i$ then determine the corner solutions and the tax-free versus taxed selling segments, exactly as in the CGT KKT system in (57), with P_t replaced by the corresponding effective trading prices.

The following algorithm ensures that all the optimality conditions summarized by the KKT conditions are satisfied:

1. Start in the non-selling region. Assume the agent is not selling at t (i.e., we are in the "non-selling" region where the sell constraint is slack). Compute the buy-side candidate by imposing the buy-side Euler equality, which prices the marginal unit at the effective purchase price $(1+\tau_t^F)P_t$:

$$C_0^i = \left(\frac{\psi_t^F}{1+\tau_t^F}\right)^{-1/\gamma}$$

Given C_0^i, use the budget constraint to obtain the associated candidate holdings at the buy price,

$$S_0^i = \frac{X_t^{F,i} - C_0^i}{\left(1+\tau_t^F\right)P_t}$$

where $X_t^{F,i} \equiv W_t + \big((1-\tau_t^D)D_t + P_t\big)S_{t-1}^i + T_t^i + \tau_t^F P_t S_{t-1}^i$. Here $X_t^{F,i}$ collects terms in the budget constraint that do not depend on S_t^i, after rewriting $-\tau_t^F|S_{t-1}^i - S_t^i|P_t = -\tau_t^F(S_t^i - S_{t-1}^i)P_t$ in the buy region. This (C_0^i, S_0^i) pair is only a buy-side candidate used to determine whether the optimum entails purchasing ($S_t^i > S_{t-1}^i$) or instead falls in the inaction (hold) region. The hold allocation itself is always $S_t^i = S_{t-1}^i$ and is characterized by the KKT inequalities (no Euler equality) rather than by this

candidate.

2. Compute the buy-side candidate by imposing the buy-side Euler equality, which prices the marginal unit at the effective purchase price $(1+\tau_t^F)P_t$: $C_0^i = \left(\frac{\psi_t^F}{1+\tau_t^F}\right)^{-1/\gamma}$ and compute the associated candidate holdings from the budget constraint using the effective buy price $S_0^i = \frac{X_t^{F,i} - C_0^i}{(1+\tau_t^F)P_t}$

3. Buy feasibility check. If the candidate indeed lies in the non-selling region, i.e. $S_0^i \geq S_{t-1}^i$, then accept $(C_t^i, S_t^i) = (C_0^i, S_0^i)$ and proceed to Step 6.

4. Otherwise, move to the selling side. Since $S_0^i < S_{t-1}^i$, consider sell-side candidates. Compute two candidate consumptions (and implied candidate share positions) using the sell-side FOC, depending on whether the marginal sale is CGT-taxed.

   (a) Tax-free sell candidate. Impose the sell-side FOC with effective marginal proceeds factor $(1-\tau_t^F)$ (i.e., no CGT on the marginal sale) and compute

   $$C_1^i = \left(\frac{\psi_t^F}{1-\tau_t^F}\right)^{-1/\gamma}$$

   This candidate is feasible in the tax-free sell regime if either $P_t \leq P_{bt}^i$ (no embedded gains) or, when $P_t > P_{bt}^i$, it lies in the loss-offset region, i.e. if $S_1^i \geq \hat{S}_t^i$. Compute $S_1^i = \frac{X_t^{S,i} - C_1^i}{(1-\tau_t^F)P_t}$, with $X_t^{S,i} \equiv W_t + \big((1-\tau_t^D)D_t + P_t\big)S_{t-1}^i + T_t^i - \tau_t^F P_t S_{t-1}^i$ (using $-\tau_t^F|S_{t-1}^i - S_t^i|P_t = -\tau_t^F(S_{t-1}^i - S_t^i)P_t$ in the selling region).

   (b) Taxed sell candidate. Impose the sell-side FOC with the effective marginal proceeds factor for a taxed marginal sale and compute

   $$C_2^i = \left(\frac{\psi_t^F}{(1-\tau_t^F) - \tau_t^{CG}\left(1 - \frac{P_{bt}^i}{P_t}\right)}\right)^{-1/\gamma}$$

   where κ_t^{CG} denotes the marginal capital-gains proceeds factor in the taxed regime (as defined above). This candidate is relevant only when $P_t > P_{bt}^i$ and $S_2^i < \hat{S}_t^i$. Using the budget constraint in the taxed-selling region, compute the implied holdings as

   $$S_2^i = \frac{X_t^{S,i} - \tau_t^{CG}\Big((P_t - P_{bt}^i) - L_t^i\Big)S_{t-1}^i - C_2^i}{P_t(1-\tau_t^F) - \tau_t^{CG}(P_t - P_{bt}^i)}$$

5. Selling feasibility check (kink at $\hat{S}_t^i$).

   (a) If the relevlant sell candidate satisfies the region condition (tax-free case: $S_1^i \geq$

$\hat{S}_t^i$; taxed case: $S_2^i \leq \hat{S}_t^i$), accept it:[59]

$$(C_t^i, S_t^i) = \begin{cases} (C_1^i, S_1^i), & \text{if tax-free sale applies,} \\ (C_2^i, S_2^i), & \text{if taxed sale applies.} \end{cases}$$

(b) Otherwise, the sell-side equality lies on the wrong side of the kink. Then the optimum is the kink $(C_t^i, S_t^i) = (X_t^{S,i} - (1-\tau_t^F)P_t\,\hat{S}_t^i, \hat{S}_t^i)$.

6. Bounds projection.

(a) If $S_t^i > \bar{S}$, set $S_t^i = \bar{S}$ and compute C_t^i from the budget constraint in the buy/hold region:

$$C_t^i = W_t + \big((1-\tau_t^D)D_t + P_t\big)S_{t-1}^i - (1+\tau_t^F)P_t\bar{S} + \tau_t^F P_t S_{t-1}^i + T_t^i$$

(b) If $S_t^i < 0$, set $S_t^i = 0$ and compute C_t^i from the budget constraint in the selling region:

$$C_t^i = W_t + \big((1-\tau_t^D)D_t + P_t\big)S_{t-1}^i - \tau_t^F P_t S_{t-1}^i - \tau_t^{CG}\max\Big\{S_{t-1}^i(P_t - P_{bt}^i) - L_t^i S_{t-1}^i, 0\Big\} + T_t^i$$

Adding a tax on unrealized gains on top of the tax on realized gains simply requires combining the baseline algorithm with the one for accrual-based taxation.

E.3. Approximating the unknown subjective conditional expectation

This section provides details on the application of the Parameterized Expectations Algorithm (PEA) used to solve the model. In Algorithm A.8, the conditional expectation $\psi(z_t^i)$ enters as an input. In equilibrium, however, $\psi(\cdot)$ is an unknown function and PEA requires specifying a parametric approximation $\mathcal{E}(z_t^i; \boldsymbol{\chi}^i)$ and estimating the parameters $\boldsymbol{\chi}^i$. Many flexible classes can serve this role (e.g., polynomials, neural networks). Here I adopt a complementary approach: I use basic economic structure to guide the approximation architecture. This parameterization is an approximation family, not an equilibrium restriction: the coefficients $\boldsymbol{\chi}^i$ are identified by the PEA loss function. In particular, let[60]

$$g_t^i \;\equiv\; \frac{\max\{P_t/D_t - P_{bt}^i/D_t - L_t^i/D_t, 0\}}{P_t/D_t}$$

[59] First, I compute the CGT-free candidate. If it lies weakly above $\hat{S}_t^i$, the marginal sale is CGT-free and this candidate satisfies the KKT system; I therefore accept it. Only if it lies below $\hat{S}_t^i$, I compute the taxed candidate; if that candidate lies weakly below $\hat{S}_t^i$, I accept it, otherwise the optimum is the kink.

[60] To render the problem stationary, I work in dividend units by normalizing nominal variables by current dividends (e.g. W_t/D_t and P_t/D_t).

which measures net embedded gains (after offsetting carryover losses) as a fraction of the price in dividend units. I approximate $\psi(z_t^i)$ by

$$\mathcal{E}(z_t^i;\boldsymbol{\chi}^i) = (\tilde{\psi}_t^i)^{-\gamma}\Big(1-\tau_t^{CG}\,x_t^i\,g_t^i\Big)^{-\gamma} \tag{231}$$

where

$$\tilde{\psi}_t^i = c_{yt}^i\Big(\frac{W_t}{D_t}+(1-\tau_t^D)S_{t-1}^i\Big)+c_{wt}^i\frac{P_t}{D_t}S_{t-1}^i$$

$$c_{yt}^i \equiv 1-\chi_0^i\delta(1-\tau_t^D)(\beta_t^d)^i$$

$$c_{wt}^i \equiv 1-\chi_0^i\delta\big(1-x_t^i g_t^i\tau_t^{CG}\big)\beta_t^i$$

and

$$x_t^i = \Lambda\big(R_t-\chi_1^i\tau_t^{CG}g_t^i\big)$$

with $\Lambda(u)=1/(1+e^{-u})$ and $R_t=(P_t+D_t)/P_{t-1}$.

This functional form is designed to be economically interpretable. The term $\tilde{\psi}_t^i$ aggregates period-t resources into a single index with linear loadings on income (including dividends) and wealth, consistent with standard linear consumption rules (e.g., Hakansson (1970), Benhabib, Cui, and Miao (2024)) and with empirical evidence on distinct propensities to consume out of income and wealth di2020stock. The factor $\big(1-\tau_t^{CG}x_t^i g_t^i\big)^{-\gamma}$ captures how expected capital-gains tax liabilities rescale marginal utility through the realization channel.

The realization nonlinearity implied by capital-gains taxation enters $\mathcal{E}$ through the kink in g_t^i, which activates the tax wedge only when the position carries net embedded gains, and through x_t^i, a smooth proxy for the probability of operating on the realization margin in that region. In this sense, x_t^i can be viewed as a reduced-form smoothed indicator for whether the household is close to the realization boundary: it is increasing in R_t and, for $g_t^i>0$, decreasing in τ_t^{CG} (consistent, for instance, with the comparative statics in Ehling, Tompaidis, and C. Yang (2019)). While extreme combinations of χ_0^i and beliefs could in principle render some coefficients negative, this does not arise in the estimated model: across all simulated paths and states used for estimation and policy counterfactuals, $\min_{t,i}\{c_{yt}^i, c_{wt}^i\}\geq 0$ and $\min_{t,i}\{1-\tau_t^{CG}x_t^i g_t^i\}\geq 0$. Thus, the approximation remains economically and numerically well behaved over the relevant region of the state space.

χ^i is an agent-specific unknown parameter vector to be estimated. For that, I follow a standard PEA procedure, except in the context of Internal Rationality. Here is the full algorithm.

1. Fix the numerical values of the structural parameters.

2. Simulate long series of the exogenous processes from the point of view of agent i $\{W_t, D_t, T_t^i, \tau_t^D, \tau_t^{CG}, P_t^i, (\beta_t^d)^i, \beta_t^i\}_{t=0}^T$ using her subjective probability measure $\mathcal{P}^i$, which involves equations (37), (38), (42), (43), (44), (45), (46), (47). The subjective

model of prices mimics equation (46). In simulations, I use a projection facility for price beliefs to avoid explosive price trajectories (see below) and augment the measurement equation for price growth with a small valuation anchor:

$$\ln \frac{P^i_{t-1}}{P^i_{t-2}} = \ln \beta^i_{t-1} - \lambda_\beta \Big(\ln(P^i_{t-1}/D_{t-1}) - \ln \bar{PD} \Big) + \varepsilon^{P,i}_{t-1}$$

where $\lambda_\beta > 0$. Thus, when the price–dividend ratio is above its long-run anchor $\bar{PD}$, agents attribute part of the realized price growth to a transitory valuation component and infer lower latent expected growth. The constant-gain rule in (232) follows by updating $\ln \beta^i_{t-1}$ using the associated forecast error. This implies the following constant-gain updating rule for perceived price-growth beliefs:[61]

$$\ln \beta^i_t = \ln \beta^i_{t-1} + g^i \left(\ln \frac{P^i_{t-1}}{P^i_{t-2}} - \ln \beta^i_{t-1} + \lambda_\beta \Big(\ln \frac{P^i_{t-1}}{D_{t-1}} - \ln \bar{PD} \Big) \right) \tag{232}$$

Note that prices are agent-specific, as they are drawn from the subjective model of prices, which is idiosyncratic. The series are kept fixed throughout the algorithm.

3. Set an initial value for the PEA parameter χ^i_0, and a starting condition for S^i_{t-1}.

4. Simulate the economy. In particular, compute allocations, the basis price, carryover losses, and the interior of the conditional expectation

$$y^i_{t+1} = \delta \Big(\frac{C^i_{t+1}}{D_{t+1}} \Big)^{-\gamma} \Big(\frac{D_{t+1}}{D_t} \Big)^{1-\gamma} \frac{D_t}{P^i_t} \Big(\frac{P^i_{t+1}}{D_{t+1}} + (1 - \tau^D_{t+1}) - \tau^{CG}_{t+1} \big(\max\{0, \frac{P^i_{t+1}}{D_{t+1}} - \frac{P^i_{b,t+1}}{D_{t+1}} \} \mathbb{1}\{S^i_{t+1} < \hat{S}^i_{t+1}\} \big) \Big) + \delta \Big(\frac{D_{t+1}}{D_t} \Big)^{1-\gamma} \frac{D_t}{P^i_t} \Big(\frac{F^i_{1,t+1}}{D_{t+1}} + \frac{F^i_{2,t+1}}{D_{t+1}} \Big) \tag{233}$$

Recall from Section E.1. that $F^i_{1,t+1}$ and $F^i_{2,t+1}$ are already expressed in marginal-utility units so that $F^i_{k,t+1}/D_{t+1}$ is dimensionless and stationary under the same dividend normalization used for prices and consumption.

5. Update χ^i. By definition of a conditional expectation,

$$y^i_{t+1} = \mathcal{E}(z^i_t; \chi) + u^i_{t+1} \tag{234}$$

[61] The expectations updating formulas (47)-(49) include lognormal (Jensen) corrections that depend on perceived conditional variances. In the simulations, I impose a calibration in which the perceived innovation variance is small relative to the perceived payoff/price-growth variance, $(\sigma^2_\nu)^i \ll (\sigma^2_d)^i$, and $(\sigma^2_p)^i = (\sigma^2_d)^i$ for simplicity. Under this restriction, the corresponding variance correction terms κ^i_d and κ^i_p are quantitatively negligible. Accordingly, I drop κ^i_d and κ^i_p from the simulated belief-updating equations and treat them as a level-normalization that does not materially affect the mechanisms operating through time variation in beliefs.

where u^i_{t+1} is the forecast error. Minimize the mean squared error to get an estimate of χ, that is

$$\hat{\chi} = \arg\min_{\chi} \sum_{t=0}^{T-1} (y_{t+1} - \mathcal{E}(z^i_t; \chi))^2 \tag{235}$$

Then, use an updating rule:

$$\chi^i_{n+1} = \lambda\hat{\chi} + (1-\lambda)\chi^i_n \tag{236}$$

where $\lambda \in (0,1)$ determines the relative weight on the new estimate. Convergence is declared when the post-update residual is small,

$$\max_i \|\hat{\chi} - \chi^i_{n+1}\|_\infty < 10^{-4}$$

Since $\hat{\chi} - \chi^i_{n+1} = (1-\lambda)(\hat{\chi} - \chi^i_n)$, this criterion is equivalent to

$$\max_i \|\hat{\chi} - \chi^i_n\|_\infty < \frac{10^{-4}}{1-\lambda}$$

so the tolerance is well-defined given $\lambda < 1$.

$\mathcal{E}(z^i_t; \chi^i)$ has a parsimonious functional form, containing only two free parameters—in sharp contrast to flexible approximators like neural networks. To assess approximation quality, I compute the one-step-ahead forecast error u^i_{t+1} for the conditional expectation object targeted by the parametric approximator. This is not an Euler residual (the Euler condition is imposed when computing allocations); rather, it measures the discrepancy between the approximated and true conditional expectation. Because both u^i_{t+1} and the consumption–dividend ratio C^i_t/D_t are dimensionless (ratios), the statistic is unit-consistent. I scale the absolute error by the contemporaneous C^i_t/D_t to obtain a magnitude normalization that is comparable across agents and time. Formally, define the consumption-scaled mean absolute error

$$\bar{\mathcal{A}}^i(\chi) = \log_{10}\left(\mathbb{E}\left[\left|\frac{u^i_{t+1}}{C^i_t/D_t}\right|\right]\right) \tag{237}$$

where $\mathbb{E}[\cdot]$ is taken over the simulated ergodic distribution (or over the relevant conditioning set, when reported for subregions of the state space). The normalization by C^i_t/D_t should therefore be read as a convenient scale adjustment—a "consumption-scaled" mean absolute error.

The histogram below shows the cross-sectional distribution of our consumption-scaled forecast-error statistic across the 100 agents. The mean is -3.86 and the median is -3.84, so the typical agent's one-step-ahead forecast error is about $10^{-3.85} \approx 1.4 \times 10^{-4}$ in consumption-scaled units—that is, roughly one part in ten thousand relative to the level

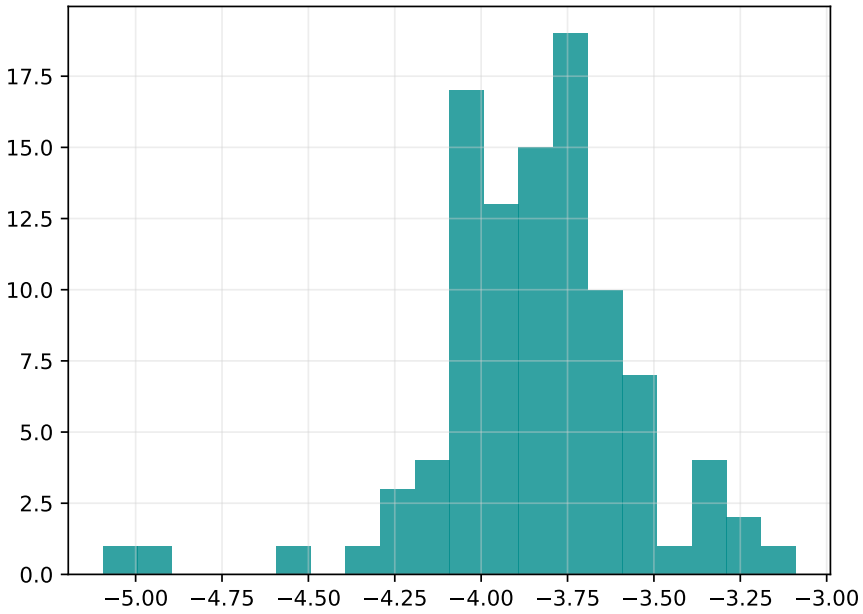


Figure 8: *Cross-sectional distribution of the accuracy statistic.*

of consumption. Economically, this indicates that the parametric expectation function closely matches the conditional mean of the pricing object that governs portfolio choices: the approximation error is tiny compared with the scale of resources that agents allocate each period.

Table 9 complements the histogram by reporting accuracy directly in levels and by conditioning on regions of the state space associated with the realization margin. The table is constructed in two steps. First, for each agent i, I compute within-agent time-series moments of the consumption-scaled absolute forecast error $e^i_{t+1} \equiv |u^i_{t+1}|/(C^i_t/D_t)$: the time-series mean, median, 90th percentile, 99th percentile, and maximum of $\{e^i_{t+1}\}_t$. Second, for each of these within-agent moments, I report cross-sectional summaries across agents (in columns). Hence, rows index moments computed over time within each agent, while columns summarize the cross-agent distribution of those time-series moments.

Consistent with the histogram, the typical agent's mean absolute error is on the order of 10^{-4}, and even tail measures (e.g., the time-series 99th-percentile error) remain well below 10^{-2}. Moreover, there is no evidence that approximation performance deteriorates in the region most relevant for inaction boundaries: conditioning on periods in which the realization margin is active and the agent sells, the corresponding cross-sectional means are comparable to—and in fact smaller than—their unconditional counterparts. Economically, these diagnostics indicate that the parametric expectation function closely tracks the conditional mean of the pricing object that governs portfolio choice, including in the subset of states that is most relevant for trading decisions.

Table 9: Accuracy of the parametric expectation approximation. *For each agent i, rows refers to within-agent time-series moments of $e^i_{t+1} \equiv |u^i_{t+1}|/(C^i_t/D_t)$. Columns report cross-sectional summaries across the 100 agents of the corresponding within-agent moment. "Global" uses the full simulated path. "Selling region" conditions the within-agent time-series moments on dates with $g^i_t > 0$ and $S^i_t > S^i_{t+1}$.*

	Across agents			
Within-agent	Mean	Median	90th pct.	99th pct.
Global (all states):				
Mean	1.76×10^{-4}	1.48×10^{-4}	3.06×10^{-4}	5.19×10^{-4}
Median	9.08×10^{-5}	8.56×10^{-5}	1.65×10^{-4}	2.38×10^{-4}
90th pct.	3.89×10^{-4}	3.25×10^{-4}	6.61×10^{-4}	1.13×10^{-3}
99th pct.	1.40×10^{-3}	9.57×10^{-4}	2.36×10^{-3}	5.95×10^{-3}
Max	4.79×10^{-3}	2.90×10^{-3}	7.76×10^{-3}	3.01×10^{-2}
Selling region (active & selling): conditional on $g^i_t > 0$ and $S^i_t > S^i_{t+1}$				
Mean	1.14×10^{-4}	1.14×10^{-4}	1.86×10^{-4}	2.20×10^{-4}
90th pct.	2.63×10^{-4}	2.56×10^{-4}	4.04×10^{-4}	5.02×10^{-4}
99th pct.	6.92×10^{-4}	6.64×10^{-4}	9.97×10^{-4}	1.21×10^{-3}
Max	1.22×10^{-3}	1.18×10^{-3}	1.78×10^{-3}	2.52×10^{-3}

Practical implementation and numerical safeguards

The baseline algorithm above describes a standard PEA update. In practice, a literal joint update of $\chi^i = (\chi^i_0, \chi^i_1)$ is numerically challenging for two reasons. First, the least-squares criterion in step 5 often exhibits weak curvature in χ^i_1 along the simulated path. Intuitively, χ^i_1 affects $\psi(\cdot)$ primarily through

$$\omega^i_t \equiv \tau^{CG}_t \frac{\max\{P_t - P^i_{bt} - L^i_t, 0\}}{P_t}$$

which is a function of accrued gains and is well defined regardless of whether the agent sells in period t. When ω^i_t is close to zero for long stretches of the sample, perturbations in χ^i_1 have little effect on $\mathcal{E}(z^i_t; \chi^i)$ and therefore little effect on the PEA regression objective. To sharpen identification, I use a weighted least-squares objective that upweights observations in which the capital-gains wedge is quantitatively important:

$$w_t \equiv \frac{\omega^{i\,2}_t}{\mathbb{E}[\omega^{i\,2}_t]}$$

so that observations with large embedded gains (and hence a large potential tax wedge) receive more weight in the PEA regression step. Nonetheless, for many agents the joint criterion remains weakly informative about χ^i_1 because χ^i_0 and χ^i_1 both tilt the fitted conditional expectation: re-optimizing χ^i_0 can partially offset changes in χ^i_1, so the profiled objective is nearly flat in χ^i_1 over economically plausible neighborhoods.

Second, the model is highly nonlinear, with kinks and occasionally binding constraints. As a result, small changes in χ_1^i may move the economy to a different regime for a subset of dates, which makes the objective non-smooth in χ_1^i and further complicates gradient-based joint optimization.

To address these issues while staying within the logic of PEA, I implement a block (two-step) update of χ^i combined with warm starts and a conservative switching rule. Conceptually, the most direct implementation would iterate on the pair (χ_0^i, χ_1^i) until convergence, i.e., alternate between solving for χ_0^i conditional on χ_1^i (the PEA regression step) and updating χ_1^i conditional on χ_0^i, repeating these two steps until the change in (χ_0^i, χ_1^i) is below tolerance. In my environment, each outer iteration requires re-simulating a highly nonlinear economy, which makes such a full outer fixed-point loop computationally costly. The procedure below is therefore designed as a computationally efficient approximation to this outer loop: I perform a local update of χ_1^i around a baseline value and then re-solve χ_0^i at the resulting $\hat{\chi}_1^i$.

Given an initial value $\chi_{1,0}^i$ for the realization-sensitivity coefficient, I proceed as follows:

1. Solve for χ_0^i conditional on χ_1^i. Holding $\chi_1^i = \chi_{1,0}^i$ fixed, iterate the PEA regression update for χ_0^i (step 5 in the baseline algorithm) until the stopping rule is satisfied. Denote the resulting coefficient by $\hat{\chi}_0^i(\chi_{1,0}^i)$. In this step, the least-squares criterion uses the weights w_t defined above.

2. Local search for χ_1 and conservative acceptance. Holding $\chi_0^i = \hat{\chi}_0^i(\chi_{1,0}^i)$ fixed, evaluate the accuracy statistic $\bar{\mathcal{A}}^i(\chi_0^i, \chi_1^i)$ over a grid of candidate values $\chi_1^i \in \mathcal{N}(\chi_{1,0}^i)$, where $\mathcal{N}(\chi_{1,0}^i)$ is a neighborhood around the initial guess. Then set

$$\hat{\chi}_1^i \in \arg\min_{\chi_1^i \in \mathcal{N}(\chi_{1,0}^i)} \bar{\mathcal{A}}^i\big(\hat{\chi}_0^i(\chi_{1,0}^i), \chi_1^i\big) \tag{238}$$

   subject to a minimum-improvement rule: adopt $\hat{\chi}_1^i \neq \chi_{1,0}^i$ only if it reduces the accuracy statistic by at least Δ,

$$\bar{\mathcal{A}}^i\big(\hat{\chi}_0^i(\chi_{1,0}^i), \hat{\chi}_1^i\big) \leq \bar{\mathcal{A}}^i\big(\hat{\chi}_0^i(\chi_{1,0}^i), \chi_{1,0}^i\big) - \Delta \tag{239}$$

   If this condition fails, set $\hat{\chi}_1^i = \chi_{1,0}^i$. This switching rule prevents mechanical selection of boundary values of χ_1^i in cases where $\bar{\mathcal{A}}^i$ is nearly monotone in χ_1^i over $\mathcal{N}(\chi_{1,0}^i)$ and improvements are economically immaterial. In practice, I set $\Delta = 0.05$. The minimum-improvement threshold Δ can be interpreted as an adjustment cost for changing χ_1^i, and it is helpful in this model with kinks and regime switching, where small numerical differences in $\bar{\mathcal{A}}^i$ may otherwise lead to large discrete changes in endogenous realizations.

3. Re-solve χ_0 at $\hat{\chi}_1$. Finally, to ensure internal coherence of the reported coefficients,

rerun the PEA regression update for χ_0^i holding $\chi_1^i = \hat{\chi}_1^i$ fixed, warm-starting from $\hat{\chi}_0^i(\chi_{1,0}^i)$. This yields the final coefficient pair $(\hat{\chi}_0^i, \hat{\chi}_1^i)$ with $\hat{\chi}_0^i = \hat{\chi}_0^i(\hat{\chi}_1^i)$ by construction.

E.4. The PEA-SMM loop

The following algorithm describes the implementation of the double PEA-SMM loop.

1. Initialization. Choose an initial guess for the structural parameters, denoted $\bar{\theta}^{(k)}$, a damping parameter $\lambda_\theta \in (0, 1]$, and a tolerance $\varepsilon_\theta > 0$.

2. Inner loop (PEA at fixed θ). Given θ, run PEA until convergence and obtain $\widehat{\chi}^i(\theta)$ for each agent i. Using these coefficients, simulate the model and compute the model-implied moment vector
$$\widetilde{\mathcal{S}}(\theta) \;\equiv\; \widetilde{\mathcal{S}}\Big(\theta; \{\widehat{\chi}^i(\theta)\}_i\Big)$$
That is, $\widetilde{\mathcal{S}}(\theta)$ is evaluated after solving the PEA problem at θ.

3. Outer loop (SMM). Estimate θ by
$$\theta^\star = \arg\min_{\theta\in\Theta} \Big[\hat{\mathcal{S}} - \widetilde{\mathcal{S}}(\theta)\Big]' \hat{W} \Big[\hat{\mathcal{S}} - \widetilde{\mathcal{S}}(\theta)\Big] \tag{240}$$
where each evaluation of the objective at a trial θ calls the inner PEA solver in Step 2.

4. Convergence check and update. If
$$\big\|\widetilde{\mathcal{S}}(\theta^\star) - \widetilde{\mathcal{S}}(\bar{\theta}^{(k)})\big\| < \varepsilon_\theta$$
stop. Otherwise, update
$$\bar{\theta}^{(k+1)} = (1 - \lambda_\theta)\bar{\theta}^{(k)} + \lambda_\theta\, \theta^\star$$
and repeat Steps 2–4 until convergence.

Acceleration via warm starts and a two-stage gate for inner PEA re-solves. Running a full inner PEA solve at every candidate parameter vector explored by the outer optimizer can be costly. I therefore implement two standard accelerations.

(i) *Warm start.* Whenever an inner PEA solve is performed at a candidate parameter vector, the fixed-point iterations are initialized at the most recently converged coefficients, $\chi^{i,(0)} \leftarrow \widehat{\chi}^i(f_{\text{last}})$, which substantially reduces inner-iteration counts.

(ii) *Two-stage evaluation with a re-solve gate.* At each outer iteration, given the current parameter vector θ, I first solve the inner PEA once to obtain agent-level coefficients

$\{\widehat{\chi}^i(\theta)\}_i$ (step 2). I then hold these inner objects fixed and let the outer optimizer search over trial parameters θ' using the corresponding moment mapping

$$\widetilde{\mathcal{S}}^{\text{scr}}(\theta' \mid \theta) \;\equiv\; \widetilde{\mathcal{S}}\Big(\theta'; \{\widehat{\chi}^i(\theta)\}_i\Big)$$

which is less expensive because it avoids re-solving the inner PEA at every function evaluation. This produces a candidate update θ_{new}. As a first screen, I require that the moments implied by θ_{new} improves relative to moments using θ.

To ensure that the candidate step remains beneficial once the inner problem is updated, I then apply a gate: I re-solve the inner PEA at θ_{new}, warm-starting from the last coefficients, to obtain $\{\widehat{\chi}^i(\theta_{\text{new}})\}_i$ and the corresponding definitive moments $\widetilde{\mathcal{S}}(\theta_{\text{new}})$. The candidate is accepted only if the moment improvement persists after this inner refresh. Upon passing the gate, I update θ using a damped step and project back to the admissible bounds.

This procedure yields the bulk of the computational gains because the outer optimizer evaluates the objective many times holding the inner solution fixed, while a full inner refresh is performed only once per outer iteration (to validate and accept the proposed step).

E.5. Projection facility

In simulations, adaptive belief updates can occasionally imply extreme valuation jumps that lie outside the region where the policy rules and conditional-expectation approximations are reliable. Following the projection-facility approach common in adaptive learning, I therefore restrict attention to an admissible set $\mathcal{PD} = (0, \overline{pd}]$ with $\overline{pd} = 275$, which is the upper bound of the price-dividend support in the subjective models of prices. Concretely, let excess demand be

$$g_t(pd; \boldsymbol{\beta}_t) \equiv \sum_i \mu^i S(z_t^i) - S_t^s$$

where S_t^s is the aggregate supply (including supply shocks). Let $\widehat{pd}_t(\beta_t)$ as a root of $g_t(\cdot; \beta_t) = 0$ preliminary equilibrium price. If the preliminary root satisfies $\widehat{pd}_t(\beta_t) \leq \overline{pd}$, accept it. Otherwise, I revert beliefs once to the previous coefficients, $\beta_t \leftarrow \beta_{t-1}$, recompute the clearing root, and set $pd_t = \widehat{pd}_t(\beta_{t-1})$. This device preserves market clearing (the accepted pd_t always solves $g_t(pd_t; \widetilde{\beta}_t) = 0$) while preventing transient belief errors from pushing the simulation into explosive or numerically unreliable regions.

E.6. An analytical result about the stock demand function

This appendix derives an as-if stock-demand representation in a representative-agent Lucas-tree economy with proportional taxes on dividends and on accrued capital gains. The exercise is intentionally narrow. Because the quantitative model features heterogeneous

agents, learning, and realization-based capital-gains taxation, there is no feasible mapping from this benchmark to the main model, and I do not use it to discipline parameter signs or comparative statics at all. Instead, the appendix serves only to illustrate that, in a tractable environment, closed-form stock demand naturally takes a "propensity × component" form with a multiplicative tax wedge, separating dividend-related from price-related terms. This motivates the functional forms adopted in Appendix E.3.

Suppose dividends follow a log random walk with drift,

$$\ln D_t = \ln \beta^D + \ln D_{t-1} + \ln \epsilon_t^d; \qquad \ln \epsilon_t^d \sim \mathcal{N}\big(-\tfrac{1}{2}s_d^2,\ s_d^2\big) \text{ i.i.d.} \tag{241}$$

so that $g_{t+1} \equiv D_{t+1}/D_t = \beta^D \epsilon_{t+1}^d$ satisfies $\mathbb{E}[g_{t+1}] = \beta^D$. The household has CRRA utility with coefficient γ and discount factor δ, and trades shares S_t of the Lucas tree at price P_t. Dividend income is taxed at rate τ^D, and accrued capital gains $(P_t - P_{t-1})S_{t-1}$ are taxed at rate τ. Define the risk moments $\Gamma \equiv \mathbb{E}[g^{\,1-\gamma}]$ and $\kappa \equiv \mathbb{E}[g^{-\gamma}]$, which are finite under (241).

The period budget constraint is

$$C_t + P_t S_t = (1-\tau^D) D_t S_{t-1} + P_t S_{t-1} - \tau\,(P_t - P_{t-1}) S_{t-1} + T_t \tag{242}$$

The government rebates contemporaneous tax revenues lump-sum,

$$T_t = \tau^D D_t S_{t-1} + \tau (P_t - P_{t-1}) S_{t-1} \tag{243}$$

and rebates occur before trading. Hence the consolidated budget constraint is $C_t + P_t S_t = (P_t + D_t) S_{t-1}$

In equilibrium, the representative agent holds the unit supply of the tree, $S_t \equiv 1$, and goods market clearing implies $C_t = D_t$.

With accrual taxation, the tax liability on the next-period capital gain is proportional to $P_{t+1} - P_t$. The Euler equation can then be written as

$$P_t = \mathbb{E}_t\left[\delta \left(\frac{D_{t+1}}{D_t}\right)^{-\gamma} \Big((1-\tau)P_{t+1} + (1-\tau^D) D_{t+1}\Big)\right] + \delta\tau\, \mathbb{E}_t\left[\left(\frac{D_{t+1}}{D_t}\right)^{-\gamma} P_t\right] \tag{244}$$

The last term collects the present value of the next-period accrual-tax charge, expressed in units of the time-t price base P_t. Applying the goods market clearing condition ($C_t = D_t$), iterating forward on prices and applying a standard transversality condition (that holds provided $\delta\,\Gamma(1-\tau) + \delta\,\kappa\,\tau < 1$), one gets

$$\frac{P_t}{D_t} = \frac{(1-\tau^D)\,\delta\,\Gamma}{1 - \delta\,\Gamma(1-\tau) - \delta\,\kappa\,\tau} \tag{245}$$

Although market clearing pins down $S_t \equiv 1$, it is useful to represent the equilibrium as if

it were generated by a linear stock-demand rule with propensities to invest out of dividend income and out of the wealth base P_t. Define

$$s_1 \equiv \delta\,\Gamma\,(1-\tau^D); \qquad s_2 \equiv \delta\,\Gamma\,(1-\tau)+\delta\,\kappa\,\tau$$

Consider the policy rule

$$S_t(P_t, S_{t-1}) = \left(\frac{s_1 D_t + s_2 P_t}{P_t}\right) S_{t-1} \tag{246}$$

Imposing last period's equilibrium $S_{t-1} = 1$ and stock-market clearing $S_t = 1$ yields $P_t/D_t = s_1/(1-s_2)$, which coincides with (245). This representation clarifies the roles of the tax instruments: τ^D reduces the effective propensity to reinvest dividend income (s_1), while the accrual tax affects the wealth-base propensity s_2.

F. Computing the non-taxable share

The evolution of the effective capital tax rates depends essentially on two factors: statutory rates and regulations. Legal regulations are accounted for by the NBER TaxSim rates. The important exception is the amount of capital income accruing to non-taxable units, as pension funds, IRAs or non-profit institutions. The Financial Accounts of the United States, run by the Fed, report the household share of corporate equity. Some takes that as a proxy for the taxable share of ownership, but that overestimate it given the inclusion of IRAs (see Rosenthal and Austin (2016) for a critical review of the different measures). Therefore, the goal is to get an estimate of the fraction of equities hold by households in taxable accounts. I follow Rosenthal and Austin (2016).

Table 10 reports the steps followed to compute the taxable share. Essentially, it amounts to an adjustment of the Fed's households equity share, considering IRAs, indirect holdings and so on. Here I detail the abbreviations dictionary: CE = corporate equities; HHNPI = households and nonprofit institutions; RoW = rest of the world; ETF = exchange traded fund; CEF = closed-end fund; REIT = real estate investment trust; C-CE = C corporations CE; MF = mutual funds; IRA = investment retirement accounts. The variables comes from the Federal's Reserve Financial Accounts of the United States, except for those variables whose construction is explained in the table. Besides, as in Rosenthal and Austin (2016), the stock held in self- directed IRAs is based on data from the Investment Company Institute. Calculations files are available upon request. Pre-1951, the share is interpolated from Sialm (2009).

Table 10: Taxable share estimation.

Step	Operation	Definition
	Total CE HHNPI	
1. Subtract foreign equities	- RoW x (Total CE HHNPI / Total CE All Sectors) = HHNPI domestic CE	
2. Subtract the stocks issued by the passthrough entities	- S Corporations – (ETF + CEF + REITs) x HH share of Mutual Funds = HHNPI domestic C-CE	
3. Subtract NPI holdings	- NPI domestic C-CE = HH domestic C-CE	NPI domestic C-CE = (NPI CE + MF stocks) x NNHPI [CE / (CE + MF)] (NPI CE + MF) given by the Fed after 1987 (CE + MF together). Before: NPI CE + MF = (HHNPI CE + MF) x (NPI CE + MF) 1987 / (HHNPI CE + MF) 1987
4. Subtract IRAs and 529 savings plans holdings	- IRA C-CE - 529 C-CE = Direct HH domestic C-CE	IRA C-CE = CE IRA x C-CE Fraction CE IRA = IRA Other Assets x 0.75 [a] 529 C-CE = College Savings Plans Assets x 0.5 [b] C-CE fraction = All sectors C-CE / (All sectors C-CE + ETF + CEF + REITs) All sectors C-CE = All sectors domestic CE – – S corp – ETF – CEF – REITs
5. Add Indirect Holdings of C Corporation Equity	+ Indirect HH domestic C-CE = HH Taxable CE	Indirect HH C-CE = (CE MF + (CE ETF + CE CEF) x HH share of MF) x Direct HH domestic CE / Total HHNPI CE
6. Divide by the total C corporation equity	/ All sector C-CE	
	= Taxable share	

[a] Assumed 75% of IRA other assets are stocks.

[b] Assumed 50% of the assets in college savings plans were C corporation equity.

[b] Acronyms: CE = corporate equity; C-CE = C-corporation equity; HH = households; NPI = nonprofit institutions; HHNPI =households and nonprofit institutions; RoW = rest of the world; MF = mutual funds; ETF = exchange-traded funds; CEF = closed-end funds; REITs = real estate investment trusts; IRA = Individual Retirement Account; 529 = 529 college savings plan; S corporations = pass-through corporations (Subchapter S).

B G. Estimating the Cross-Sectional Distribution of Kalman Gains

I use repeated cross-sections from the UBS/Gallup survey on investors' return expectations. Each survey wave provides individual-level expectations and demographics, but the survey does not follow individuals over time.

I follow standard data-cleaning procedures. Observations with missing data on expected returns, age, income, or sampling weights are dropped. I also exclude implausibly large responses with expected returns exceeding 95 percentage points in absolute value. These filters reduce the sample from 97,712 to 76,643 observations over the period 1998:m5–2007:m6.

Because the UBS survey is a sequence of cross-sections rather than a panel, I construct synthetic cohort panels by grouping respondents into stable demographic cells. I partition investors into four age bins (35–45, 45–55, 55–65, 65–80) and income quartiles based on the income distribution in the cleaned sample, yielding 16 cohorts.

To match the quarterly frequency of the theoretical model, I aggregate individual-level observations directly to the quarterly level. Specifically, within each cohort c and quarter t, I pool all individual survey responses from the three months of that quarter and compute a single weighted mean using UBS sampling weights w_i:

$$\bar{R}^e_{c,t} \;=\; \frac{\sum_{i\in\mathcal{I}_{c,t}} w_i\, \mathbb{E}^i[R_{t,t+4}]}{\sum_{i\in\mathcal{I}_{c,t}} w_i} \tag{247}$$

where $\mathcal{I}_{c,t}$ denotes the set of individuals in cohort c surveyed during quarter t, and $\mathbb{E}^i[R_{t,t+4}]$ is individual i's expected nominal net return over the next year. This aggregation procedure preserves the micro-structure of the data: each individual response contributes to the quarterly cohort mean according to its sampling weight, regardless of which month within the quarter the response was recorded.[62]

Following Adam et al. (2015), I use own-portfolio expected returns, which are available throughout the sample period 1998:Q2–2007:Q2 and are highly correlated with expected market returns (available only until 2003.m4).

To ensure consistency with the theoretical model, I estimate the Kalman gain using expected price growth rather than total returns.[63] I transform the cohort mean expected nominal return into expected real price growth in two steps. First, I subtract the median one-year-ahead CPI inflation forecast from the Survey of Professional Forecasters (SPF)

[62] In the model, $\beta_{c,t}$ is an end-of-quarter posterior after observing P_t/P_{t-1}. In the data, $\beta_{c,t}$ pools monthly survey waves within quarter t, so it averages beliefs formed at different intra-quarter dates based on partially realized price growth. Using full-quarter P_t/P_{t-1} in (251) is therefore an approximation.

[63] Estimating the Kalman gain using total expected returns yields nearly identical results.

to obtain expected real returns:

$$\bar{R}^{e,\text{real}}_{c,t} \;=\; \bar{R}^{e}_{c,t} - \pi^{e,\text{SPF}}_{t} \tag{248}$$

where $\pi^{e,\text{SPF}}_{t} \equiv \mathbb{E}^{\text{SPF}}_{t}[\Pi_{t,t+4}]$ is the median expected net inflation rate over the next four quarters. Second, I subtract the dividend yield component scaled by expected dividend growth to isolate expected price growth:

$$\mathbb{E}^{c}_{t}\left[\frac{P_{t+4}}{P_t} - 1\right] \;=\; \bar{R}^{e,\text{real}}_{c,t} - \beta^{D}\frac{D_t}{P_t} \tag{249}$$

where $\beta^{D} \equiv \mathbb{E}[D_{t+4}/D_t]$ is the average gross annual real dividend growth over the sample period, and D_t/P_t is the dividend-price ratio (annualized).

Finally, I convert annual expected price growth to quarterly frequency. Let $\beta^{(4)}_{c,t} \equiv \mathbb{E}^{c}_{t}[P_{t+4}/P_t]$ denote the expected gross annual price growth. The implied quarterly expectation is:

$$\beta_{c,t} \;\equiv\; \mathbb{E}^{c}_{t}\left[\frac{P_{t+1}}{P_t}\right] \;=\; \left(\beta^{(4)}_{c,t}\right)^{1/4} \tag{250}$$

This transformation assumes that investors expect price growth to be approximately constant across quarters within the year.

For each cohort c, I estimate the Kalman gain g^{c} from the belief updating equation:

$$\beta_{c,t} \;=\; \beta_{c,t-1} + g^{c}\left(\frac{P_t}{P_{t-1}} - \beta_{c,t-1}\right) + u_{c,t} \tag{251}$$

where P_t/P_{t-1} is realized quarterly gross real price growth. To estimate it, I subtract $\beta_{c,t-1}$ from both sides so that the specification regresses the revision in beliefs on the forecast error, with the Kalman gain g^{c} measuring the responsiveness of cohort c's expectations to new information. I estimate (251) by OLS without an intercept, and compute standard errors using the Newey-West HAC estimator with four lags to account for potential serial correlation in the residuals.

Table 11 reports the cohort-specific constant-gain estimates. The estimated gains are in line with the standard values in the literature, typically between 0.02 and 0.04, implying gradual updating of expected price growth in response to realized price growth innovations (see, for instance, Adam et al. (2015), Adam, Marcet, and Nicolini (2016), Adam, Marcet, and Beutel (2017)).[64]

Using the cross-section of cohort-level constant-gain estimates in Table 11, I summarize

[64] One cohort (Age 35–45, Income \$15,000–45,000) yielded a negative point estimate (−0.0006, SE = 0.0150), but this is insignificantly different from zero. It is excluded from the log-Normal estimation below.

Table 11: Cohort-specific constant-gain estimates. *Each row reports the cohort-specific constant-gain estimate $\hat{g}$ from (251). Income bin refers to quartiles of the income distribution (values shown in USD, based on category midpoints). Standard errors use Newey-West with 4 lags. The number of cohort-quarter observations used in estimation after forming $\Delta\bar{\beta}_{c,t}$ and $\bar{\beta}_{c,t-1}$ is 36.*

Age bin	Income bin	$\hat{g}$	Std. Errors
35–45	15,000–45,000	−0.0006	(0.0150)
55–65	67,500–87,500	0.0185	(0.0130)
45–55	15,000–45,000	0.0223	(0.0091)
65–80	15,000–45,000	0.0236	(0.0051)
35–45	45,000–67,500	0.0244	(0.0046)
35–45	67,500–87,500	0.0248	(0.0070)
55–65	15,000–45,000	0.0259	(0.0092)
35–45	87,500–200,000	0.0271	(0.0055)
45–55	45,000–67,500	0.0290	(0.0036)
45–55	87,500–200,000	0.0303	(0.0054)
65–80	67,500–87,500	0.0333	(0.0131)
65–80	45,000–67,500	0.0335	(0.0087)
65–80	87,500–200,000	0.0346	(0.0140)
55–65	87,500–200,000	0.0349	(0.0090)
55–65	45,000–67,500	0.0363	(0.0067)
45–55	67,500–87,500	0.0408	(0.0049)

learning-rate variation by fitting a log-normal distribution,

$$g^c \sim \log N\left(\ln(\mu_g) - \frac{1}{2}\ln\Big(1+\frac{\sigma_g^2}{\mu_g^2}\Big),\ \ln\Big(1+\frac{\sigma_g^2}{\mu_g^2}\Big)\right),$$

where $\mu_g \equiv \mathbb{E}[g^c]$ and $\sigma_g^2 \equiv \mathrm{Var}(g^c)$ denote the mean and variance of the (latent) cohort gains in levels. Since the log-normal likelihood has support on $(0,\infty)$, I estimate the distribution using the 15 cohorts with $\hat{g}^c > 0$ (the single negative point estimate, $\hat{g} = -0.0006$, is economically negligible and statistically indistinguishable from zero given its HAC standard error). Maximum-likelihood estimation yields $\hat{\mu}_g = 0.0293$ and $\hat{\sigma}_g = 0.0062$. Importantly, the dispersion in $\{\hat{g}^c\}$ should not be interpreted as direct evidence of true cross-cohort heterogeneity.

Because each cohort gain is itself estimated with nontrivial sampling uncertainty, the raw cross-cohort spread in $\hat{g}^c$ mechanically mixes heterogeneity and estimation noise. A simple measurement-error decomposition, $\hat{g}^c = g^c + u^c$ with $\mathrm{Var}(u^c) = s_c^2$, implies $\mathrm{Var}(\hat{g}^c) = \mathrm{Var}(g^c) + \mathbb{E}[s_c^2]$, so the log-normal MLE dispersion is conservative in the sense that it provides an upper bound on the underlying cross-cohort variance $\mathrm{Var}(g^c)$.