

Tax-Rate Risk: Pricing and Hedging Legislative Uncertainty in After-Tax Portfolios

Anirban Majumdar

Abstract

Almost every tax-aware calculation in wealth management assumes today's capital-gains rate persists until the investor sells. Statute says otherwise. The top federal long-term rate has moved between 15% and 28% five times since 1978, and the scheduled sunset of current law keeps the question live. Deferring a gain is therefore not a free option: it converts a tax liability known at today's rate into one owed at an unknown future rate. We show that holding an unrealized gain is economically a short position in the future tax rate, and we price it. A two-date result gives the deferral advantage in closed form, $\Delta = g_0(\tau_0 e^{\rho T} - \tau_T)$, and a certainty-equivalent haircut for rate uncertainty equal to $\frac{\gamma}{2} g_0^2 \mathrm{Var}[\tau_T]$, quadratic in the embedded gain. We then embed the realization decision in a continuous-time regime-switching model and solve the resulting optimal-stopping problem. Under constant or mean-reverting rates the investor never voluntarily realizes a gain, recovering the Constantinides tax-timing policy. Under a risk-neutral rate drift above a computable threshold the free boundary opens, and the optimal policy accelerates realizations: in our illustrative calibration to the US rate path, that threshold corresponds to an expected rate rise the investor can read off a single inequality. The practical takeaway is that the largest, lowest-basis positions carry the most rate risk, since the premium grows with the square of the embedded gain, and that the moment to accelerate realization is when a low-rate regime meets a rising-rate expectation near a forced sale.

Anirban Majumdar (InnovationStrat Consulting, LLC)

Working paper. Prepared for SSRN.

1. Introduction

A taxable investor with a large embedded gain faces a decision every year: sell now and pay tax at the current rate, or hold and pay later at whatever rate then prevails. The standard advice is to defer. Deferral lets the untaxed portion of the gain keep compounding, so the government is effectively extending an interest-free loan equal to the deferred tax. Constantinides (1983, 1984) formalized this into the tax-timing option: with symmetric treatment of gains and losses and a constant tax rate, the investor should realize losses immediately and defer gains indefinitely. The result is clean, and it underlies how direct-indexing and tax-aware long-short mandates are marketed today.

The constant-rate assumption is doing quiet work. The moment we admit that the rate at the future realization date is itself uncertain, the deferral decision stops being a one-sided bet. When you defer, you no longer know the rate at which the gain will eventually be taxed. If rates rise

between now and your sale, you have deferred into a more expensive world, and the compounding benefit has to cover that loss before deferral pays. Deferring an embedded gain g_0 is, in cash-flow terms, a commitment to pay $\tau_T \cdot g_0$ at a random future rate τ_T instead of $\tau_0 \cdot g_0$ today. That is a short position in τ_T , and short positions carry risk premia.

This paper prices that position. The contribution has three parts. First, a two-date model gives the deferral advantage and its certainty-equivalent value in closed form, isolating a *tax-rate-risk premium* that scales with the square of the embedded gain and the variance of the future rate. Second, we model the applicable rate as a continuous-time Markov chain over legislative regimes and cast the realization decision as an optimal-stopping problem; the resulting Hamilton-Jacobi-Bellman variational inequality has a free boundary that separates "keep deferring" from "realize now." Third, we characterize how that boundary moves with the risk-neutral drift of rates and give the acceleration rule: realize when the risk-neutral expected future rate rises above today's rate grossed up by the compounding factor. Numbers throughout are illustrative and simulated with a fixed seed, calibrated loosely to the historical path of US long-term capital-gains rates.

What is new relative to the tax-aware canon (Arnott, Berkin and Ye 2001; Berkin and Ye 2016; Sosner, Krasner and Liberman 2019; Israel, Moskowitz, Ross and Serban 2020) is the object of study. That literature quantifies the value of harvesting and deferral taking the rate schedule as fixed. Here the rate schedule is the source of risk. Relative to the dynamic-tax literature (Constantinides 1983, 1984; Dammon, Spatt and Zhang 2001, 2004), the novelty is that rates are stochastic, which converts a deterministic timing option into a genuine optimal-stopping problem with a rate-driven free boundary and a priced risk. A practitioner reading this should come away able to say when deferral is a bet worth making and roughly how large the wager is.

2. Related work

The tax-timing option originates with Constantinides (1983), who showed in a capital-market-equilibrium setting that the ability to time realizations has value, and Constantinides (1984), who derived the optimal trading policy: defer gains, realize losses at once. Dybvig and Koo, and later Dammon, Spatt and Zhang (2001), extended the analysis to dynamic consumption and portfolio choice with capital-gains tax; Dammon, Spatt and Zhang (2004) added asset location across taxable and tax-deferred accounts. Every one of these treats the tax rate as a known constant.

A separate empirical literature studies how realizations respond to rate *changes*. Feldstein, Slemrod and Yitzhaki (1980) documented the lock-in effect, whereby high rates discourage selling. Poterba and others have measured the elasticity of realizations to statutory rates, including the surge in selling ahead of pre-announced increases. Sialm (2009) showed that expected tax-rate changes are priced in equilibrium asset returns. That work tells us investors clearly do respond to rate risk. It does not give the portfolio manager a valuation of the option they hold or a rule for when to exercise it.

The practitioner tax-aware literature is where the deferral benefit is quantified for real portfolios. Arnott, Berkin and Ye (2001) and Berkin and Ye (2016) estimate the value of systematic loss harvesting and FIFO lot accounting; Chaudhuri, Burnham and Lo (2020) evaluate harvesting alpha out of sample; Sosner, Krasner and Liberman (2019) and Sosner, Gromb, Nemtchinov and Pyne (2021) decompose the after-tax value of a tax-aware long-short strategy into deferral,

net-loss harvesting, and character. Sialm and Sosner (2018) analyze the character advantage of the short book, and Goldberg, Cai and Sosner (2020) study the turnover-versus-tax trade-off. Convex-optimization treatments (Moehle and Boyd 2020; Haugh, Iyengar and Wang) build the harvesting problem as a dynamic program. These frameworks are the right home for the risk we price, and none of them currently carry a stochastic rate. This paper supplies the missing piece and shows where it changes the answer.

3. Model

3.1 Setup

A taxable investor holds one lot of a security. Normalize the current price to $P_0 = 1$ per unit, with cost basis $B \in [0, 1)$ and embedded gain $g_0 = 1 - B > 0$. The security's value follows $S_t = e^{\rho t}$, where ρ is the after-tax rate at which the retained gain compounds if left invested. Treating the price path as deterministic is a modeling choice, and a deliberate one: it strips out market risk so the analysis isolates *tax-rate* risk. Section 6 reintroduces price volatility and shows it does not overturn the conclusions.

The capital-gains tax rate at time t is τ_t , a stochastic process. The investor pays tax $\tau_t \cdot$ (realized gain) whenever a gain is realized. We assume symmetric treatment (losses offset gains at the same rate) so that the pure effect of rate uncertainty stands out; wash-sale and character frictions are studied in the companion papers of this series and are held aside here.

Two questions organize the analysis. When should the investor realize the embedded gain, and what is the deferral option worth once the future rate is random?

3.2 The deferral advantage in closed form

Start with a horizon T at which the position is liquidated regardless (a consumption or rebalancing need). Compare two policies for the *embedded* gain g_0 .

Defer. Hold to T and realize. Terminal wealth is

$$W_{\text{def}} = S_T - \tau_T(S_T - B) = P_0X - \tau_T(P_0X - B),$$

where $X = e^{\rho T}$ is the gross growth factor and τ_T is the (random) rate at T .

Realize now and repurchase. Pay $\tau_0 g_0$ today, funded from a reserve that would otherwise have grown by the same factor $X' = e^{\rho T}$, and repurchase the security so market exposure is unchanged. Basis resets to P_0 . At T the new gain $P_0(X - 1)$ is taxed at τ_T . Terminal wealth, netting the opportunity cost of the reserve, is

$$W_{\text{real}} = P_0X - \tau_T P_0(X - 1) - \tau_0 g_0 X'.$$

Proposition 1 (deferral advantage). The terminal-wealth advantage of deferring over realizing now is

$$\Delta = W_{\text{def}} - W_{\text{real}} = g_0 (\tau_0 X' - \tau_T).$$

Derivation. Subtract the two expressions:

$$\Delta = -\tau_T(P_0X - B) + \tau_TP_0(X - 1) + \tau_0g_0X'.$$

Write $P_0X - B = P_0(X - 1) + (P_0 - B) = P_0(X - 1) + g_0$. The two $\tau_TP_0(X - 1)$ terms cancel, leaving $-\tau_Tg_0 + \tau_0g_0X' = g_0(\tau_0X' - \tau_T)$. ■

The cancellation is the useful part. Gains accrued after today are taxed at τ_T under either policy, so they drop out, and what remains is the treatment of the *original* embedded gain: pay τ_0g_0 now and let it compound to τ_0g_0X' , or pay τ_Tg_0 at the horizon. Deferral wins exactly when

$$\tau_T < \tau_0X' \iff \frac{\tau_T}{\tau_0} < e^{\rho T}. \quad (1)$$

Under a constant rate $\tau_T = \tau_0$, condition (1) always holds and $\Delta = \tau_0g_0(X' - 1) > 0$: deferral is unambiguously good, which is the classical result. Rate risk enters through τ_T alone, and it enters with a minus sign. The derivative $\partial\Delta/\partial\tau_T = -g_0$ says the deferring investor is short g_0 units of the future rate.

3.3 The tax-rate-risk premium

Take expectations over τ_T with X' deterministic:

$$\mathbb{E}[\Delta] = g_0(\tau_0e^{\rho T} - \mathbb{E}[\tau_T]), \quad \text{Var}[\Delta] = g_0^2 \text{Var}[\tau_T]. \quad (2)$$

Deferral pays in expectation when $\mathbb{E}[\tau_T] < \tau_0e^{\rho T}$, and its variance is the rate variance scaled by the square of the embedded gain. For an investor with mean-variance preferences (or CARA utility under an approximately normal τ_T) and risk aversion γ , the certainty-equivalent value of the deferral bet is

$$\text{CE}[\Delta] = \mathbb{E}[\Delta] - \frac{\gamma}{2} \text{Var}[\Delta] = g_0(\tau_0e^{\rho T} - \mathbb{E}[\tau_T]) - \frac{\gamma}{2} g_0^2 \text{Var}[\tau_T]. \quad (3)$$

The last term is the **tax-rate-risk premium**, the haircut a risk-averse investor applies to the headline deferral benefit for bearing legislative uncertainty:

$$\Pi_{\text{rate}} = \frac{\gamma}{2} g_0^2 \text{Var}[\tau_T]. \quad (4)$$

It is quadratic in the embedded gain. Doubling the unrealized gain quadruples the rate-risk premium, which is why the largest low-basis positions, the ones deferral is supposed to help most, carry the most rate risk. It also grows with the variance of the future rate, so an environment of contested tax policy is materially different from a settled one even when the expected rate is the same.

If tax-rate risk is priced in the wider economy (rates tend to rise in fiscally stressed states that coincide with high marginal utility), the relevant expectation for the *decision* is the risk-neutral one. Replacing $\mathbb{E}$ by $\mathbb{E}^{\mathbb{Q}}$ in (2) gives the acceleration rule in its sharpest form:

$$\text{realize now if } \mathbb{E}^{\mathbb{Q}}[\tau_T] \geq \tau_0 e^{\rho T}. \quad (5)$$

3.4 The realization-timing option under regime-switching rates

The two-date model answers "defer or not" for a fixed horizon. The richer question is *when* to realize when the investor may act at any time and rates evolve continuously. Let the tax rate follow a continuous-time Markov chain $\eta_t \in \{1, \dots, K\}$ with generator $Q = (q_{ij})$ and rate levels $\tau^{(1)} < \dots < \tau^{(K)}$, so $\tau_t = \tau^{(\eta_t)}$. Legislative changes are the jumps of η_t . The investor holds the lot to a forced horizon $\bar{T}$ and may realize the embedded gain once at a stopping time $\theta \leq \bar{T}$, after which the position is repurchased with fresh basis. Because gains accrued after a reset are taxed identically whether or not the reset happened, the single-realization restriction captures the economically relevant decision; we relax it in Section 6.

Let $V(t, S, i)$ be the expected terminal after-tax wealth from managing the lot optimally, given time t , price S , and regime i . Stopping (realizing) yields the value of a clean, repurchased lot, $\Psi(t, S, i)$. Deferring keeps the option alive, and in the continuation region V solves the Kolmogorov/accrual equation

$$\partial_t V + \rho S \partial_S V + \sum_j q_{ij} V(t, S, j) = 0,$$

where $\rho S \partial_S V$ is the deterministic price accrual and $\sum_j q_{ij} V(\cdot, j)$ (with $q_{ii} = -\sum_{j \neq i} q_{ij}$) is the regime generator acting on V . Optimality across the two actions gives the variational inequality

$$\min \left\{ -\partial_t V - \rho S \partial_S V - (QV)(t, S, i), \quad V(t, S, i) - \Psi(t, S, i) \right\} = 0, \quad (6)$$

with terminal condition $V(\bar{T}, S, i) = S - \tau^{(i)}(S - B)$ (forced liquidation at the prevailing rate). The realize region is $\{V = \Psi\}$ and the defer region is $\{V > \Psi\}$, where the PDE in (6) holds. The boundary between them is the object of interest.

Proposition 2 (constant rates recover Constantinides). If $Q = 0$ and $\tau^{(i)} \equiv \tau$, the realize region is empty for any gain $g_0 > 0$; the investor defers at every (t, S) .

Proof. With frozen rates, stopping at t versus deferring to $\bar{T}$ is exactly the two-date comparison of Proposition 1 with $\tau_T = \tau$, so the advantage of deferring is $\tau(S_t - B)(e^{\rho(\bar{T}-t)} - 1) > 0$ everywhere. Deferral strictly dominates realization, so $V > \Psi$ throughout and the stopping set is empty. For a loss ($S_t < B$) the sign flips and immediate realization dominates, recovering the asymmetric harvest-losses-defer-gains policy. ■

Proposition 3 (positive drift opens the boundary). Suppose the generator has a positive risk-neutral drift, meaning $\mathbb{E}^{\mathbb{Q}}[\tau_{\bar{T}} \mid \eta_t = i]$ exceeds $\tau^{(i)}$ in the low regimes. Then there is a horizon proximity $\epsilon^* > 0$ such that for $\bar{T} - t < \epsilon^*$ the realize region is nonempty in the lowest regime, and ϵ^* increases in the drift.

Proof sketch. Near the horizon, $\bar{T} - t = \epsilon$ and $e^{\rho\epsilon} - 1 = \rho\epsilon + O(\epsilon^2)$, so the compounding benefit vanishes at rate ϵ . From (2) the deferral advantage in regime i is $\Delta \approx g(\tau^{(i)} e^{\rho\epsilon} - \mathbb{E}^{\mathbb{Q}}[\tau_{\bar{T}} \mid i]) = g(\tau^{(i)} - \mathbb{E}^{\mathbb{Q}}[\tau_{\bar{T}} \mid i]) + O(\epsilon)$. In a low regime with upward drift the leading term is negative, so for

small enough ϵ deferral is dominated and stopping is optimal. A larger drift makes the leading term more negative and pushes ϵ^* out, enlarging the realize region. ■

Proposition 3 is the analytical content of the paper. It says the deferral policy is not universal once rates can rise: an investor in a low-rate regime, close enough to a forced realization, should lock in the low rate rather than gamble on a legislated increase. The direction is worth pausing on. You accelerate realization in the *low* regime, not the high one, because a low current rate is the cheap opportunity that a rising-rate process is likely to take away. In a high regime you keep deferring and hope for reversion. This is the opposite of the intuition that a high current rate should make you want to sell.

3.5 Numerical scheme

Equation (6) is solved by explicit backward induction on a time grid $\{t_k = k \Delta t\}$. The exact regime transition over one step is $\Pi = e^{Q\Delta t}$, computed with a matrix exponential, so the regime dynamics carry no discretization error. The deterministic price collapses the S -dimension to $S_{t_k} = e^{\rho t_k}$, leaving a small (t, i) grid. At each node the algorithm compares the stop value $\Psi(t_k, S_{t_k}, i)$, computed as the expected terminal wealth of a repurchased lot using $\mathbb{E}[\tau_{\bar{T}} \mid \eta_{t_k} = i]$ from Π , against the continuation value ΠV_{k+1} , and takes the larger. The full routine is in `code/tax_rate_risk.py`.

4. Data and calibration

The model needs a rate process. We calibrate it loosely to the historical path of the US top statutory long-term capital-gains rate, which we treat as illustrative rather than as an estimated law of motion. The path is stored in `data/ltcg_top_rates.csv`. The top rate was near 28% in 1978-80, fell to 20% under the 1981 act, returned to 28% after the Tax Reform Act of 1986, dropped to 20% in 1997, to 15% in 2003, and settled at 23.8% from 2013 once the 3.8% net investment income tax is included. Over roughly five decades the rate has lived in a band from 15% to 28%, moved by six major acts, and reverted rather than trended.

We map this into a three-state Markov chain with rates $\{0.15, 0.238, 0.31\}$: a low regime, the current regime, and a high regime slightly above the historical peak to allow for the scheduled expiry of current law and for fiscal pressure. The base switching intensity is 0.15 per year on each link, which corresponds to about one legislative move per seven years and matches the historical cadence of major tax acts (1981, 1986, 1997, 2003, 2013, 2017). A single drift parameter $\delta \in [-1, 1]$ tilts the generator: at $\delta = 0$ the chain is symmetric and mean-reverting to the current rate; at $\delta > 0$ the upward intensities dominate and the risk-neutral expected future rate rises. The asset growth rate is $\rho = 6\%$, risk aversion $\gamma \in \{1, 3\}$, and horizons run to 30 years. All Monte-Carlo results use 400,000 paths and the fixed seed 20260804, so every number reproduces exactly. The regime states, generator, and parameters are written to `data/params.json` on each run.

I should be plain about the status of this calibration. It is a stylized rate process whose band and cadence echo the historical record, chosen to make the mechanism legible rather than to forecast anything. The "high" state is a scenario the investor might worry about, and readers with a different worry can substitute their own states. The qualitative results (the sign of the

boundary shift, the quadratic premium) do not depend on the exact numbers. The magnitudes do, and should be read as illustrative.

5. Results

5.1 The deferral benefit and how rate drift erodes it

Table 1 reports the deferral advantage $\mathbb{E}[\Delta]$ as a percentage of the embedded gain, for the constant-rate benchmark, the symmetric mean-reverting chain, and an upward-drifting chain ($\delta = 0.5$), across horizons.

Table 1. Deferral value as a percent of the embedded gain.

Scenario	T = 5y	T = 10y	T = 20y	T = 30y
Constant (Constantinides)	8.3%	19.6%	55.2%	120.2%
Mean-reverting	8.8%	20.1%	55.8%	120.7%
Trending up ($\delta = 0.5$)	5.6%	15.9%	51.0%	115.9%

Mean reversion turns out to be nearly free. A symmetric rate process leaves the deferral benefit almost unchanged from the constant-rate case, and in fact a hair higher, because the option to realize in a future low regime has small positive value. An upward drift is what costs the deferring investor. At a twenty-year horizon the benefit falls from 55.2% of the gain to 51.0%, a reduction of about four percentage points of the embedded gain, and the gap is proportionally larger at short horizons (5.6% versus 8.3% at five years). The long-horizon equity holder is well insured by compounding; the short-horizon or forced seller is exposed.

5.2 When deferral destroys value

Figure 2 makes the reversal explicit. The left panel plots the deferral value against the risk-neutral drift for horizons of 2, 5, 10, and 20 years. The 20-year curve never crosses zero within the plotted range: at long horizons the interest-free-loan benefit is so large that no plausible rate rise erases it. The 2-year curve crosses zero at a drift near 0.94, and shorter horizons cross sooner. The right panel fixes a steep upward drift ($\delta = 0.8$) and sweeps the asset's growth rate. Here the message is about which assets are exposed. A five-year holder of a low-growth asset (a bond sleeve returning below about 3.5% per year) sees deferral turn negative under a steep drift, while a 20-year equity holder stays comfortably positive. Deferral destroys value in the corner where compounding is weak and expected rate increases are strong: short horizons, low-return assets, and pre-announced hikes.

This is the honest form of the paper's warning. Deferral is not a losing bet in general. It becomes one only when the government's interest-free loan is small (short time, low return) and the future rate is expected to jump. The scheduled sunset of a low-rate regime is the textbook case, and it is the setting in which practitioners observe realization surges.

5.3 The tax-rate-risk premium

Table 2 reports the certainty-equivalent haircut $\Pi_{\text{rate}} = \frac{\gamma}{2} g_0^2 \text{Var}[\tau_T]$ from (4), in cents per dollar of position value, at a twenty-year horizon with $\gamma = 3$. Rows vary the embedded gain; columns

vary the standard deviation of the future rate.

Table 2. Tax-rate-risk premium (CE haircut, cents per \$1 of position, $T = 20y$, $\gamma = 3$).

Embedded gain g_0	$sd(\tau_T)$	0.03	0.05	0.07
0.25		0.01c	0.02c	0.05c
0.50		0.03c	0.09c	0.18c
0.75		0.08c	0.21c	0.41c
0.90		0.11c	0.30c	0.60c

The premium is small in absolute terms and sharply convex in the gain. Moving from a half-embedded gain to a nine-tenths-embedded gain at a five-point rate volatility roughly triples the haircut, from 0.09 to 0.30 cents. Against a deferral benefit worth tens of cents on the dollar (Table 1), the rate-risk premium is a second-order correction for most investors. It becomes first-order for the concentrated, very-low-basis holder in a high-rate-volatility environment, the client for whom deferral is otherwise most attractive. The right panel of Figure 3 shows the quadratic shape directly: the haircut curves upward in the embedded gain and fans out with rate volatility.

5.4 The optimal-stopping boundary

Table 3 compares the value of following the optimal stopping policy against two benchmarks, always-defer and realize-now-at-the-start, for a lot with basis 0.40 (embedded gain 0.60) starting in the current 23.8% regime.

Table 3. Optimal-policy value versus always-defer and realize-now (per \$1 of position).

Scenario	T	Optimal	Always-defer	Realize-now	Realizes early?
Constant	10	1.4837	1.4837	1.3942	no
Constant	20	2.6251	2.6251	2.3727	no
Mean-reverting	10	1.4944	1.4912	1.3979	no
Mean-reverting	20	2.6476	2.6407	2.3833	no
Trending ($\delta = 0.5$)	10	1.4376	1.4315	1.3683	no
Trending ($\delta = 0.5$)	20	2.5129	2.5027	2.2893	no
Trending ($\delta = 0.8$)	10	1.4064	1.4020	1.3537	yes
Trending ($\delta = 0.8$)	20	2.4476	2.4427	2.2484	yes

Under constant and mean-reverting rates the optimal policy coincides with always-defer, as Proposition 2 requires; the "realizes early" column is empty and realize-now leaves nine to twenty-five cents per dollar on the table. Under a steep drift ($\delta = 0.8$) the optimal policy begins to realize early, and it beats always-defer, though the dollar improvement is modest (2.4476 versus 2.4427 at twenty years, about half a cent). Searching over the drift, the boundary first opens at $\delta \approx 0.6$ for this lot at a twenty-year horizon. The timing option is real, its sign is unambiguous, and its dollar value is second-order relative to the first-order deferral benefit. That ordering is itself a result: for most investors the important decision is whether to defer at all, and only for large gains facing a steep pre-announced increase does the fine timing of realization move the needle.

Figure 1 plots the realize region in (t, regime) space for $\delta = 0.6$. The realize region sits in the low regime, from about year 12 to the horizon, with a thin sliver in the current regime just before liquidation. The high regime is always in the defer region. This is Proposition 3 made visible: realize into low rates as the horizon approaches, keep the option open when rates are already high.

6. Robustness and triangulation

Three checks support the main results.

Recovering the constant-rate limit. Setting the switching intensity to near zero collapses the regime chain to a frozen rate, and the optimal-stopping solver reproduces the analytic Constantinides policy: the realize region is empty and the optimal value equals always-defer to four decimals (the "Constant" rows of Table 3). This confirms the numerical scheme against the known closed-form limit.

Two engines, one answer. The deferral value is computed two independent ways. The closed-form Proposition 1, evaluated by Monte Carlo over simulated terminal rates, produces Table 1 and Figure 2. The backward-induction dynamic program of Section 3.5 produces Table 3. Where they overlap (the value of deferring versus realizing now), they agree in sign and ordering across all scenarios, and the DP's optimal-versus-always-defer gap is consistent with the small option value implied by the closed form near the boundary.

Adding price volatility. The deterministic-price assumption isolates rate risk but is not innocuous, so we re-ran the deferral comparison with the security following a geometric Brownian motion, $dS/S = \mu dt + \sigma dW$ with $\mu = 6\%$ and $\sigma = 18\%$, tax rate and asset return simulated jointly. Price risk widens the distribution of the deferral advantage but does not shift its mean materially or change the sign of the rate-drift effect: the erosion of deferral value under upward drift, and the quadratic-in-gain shape of the premium, both survive. Price risk and rate risk are close to independent in their effect on the deferral decision, which is what the separability of the terminal-wealth expression in (2) would predict. The single-realization restriction was also relaxed by allowing repeated stopping; because gains accrued after a reset are taxed identically, the multi-stop policy adds only the small extra value of harvesting future low-regime opportunities, and the boundary keeps the same shape.

I will note one thing that surprised me. Going in, I expected a high current rate to be a reason to sell (get out before it gets worse). The model says the opposite. A high current rate is a reason to *wait*, because a mean-reverting or even mildly drifting process is more likely to hand

you a lower rate later, and it is a low current rate that you should be rushing to use before it disappears. The lock-in intuition is about levels; the timing option is about the drift.

7. Limitations

The rate process is illustrative. We calibrate its band and cadence to the US historical path but do not estimate a generator from data, and the "high" state is a scenario. An investor with a genuine view on the risk-neutral rate drift can read the acceleration rule (5) directly; an investor without one should treat the mean-reverting case as the base and the drift cases as stress tests. The model also abstracts from the character dimension (short-term versus long-term rates can move differently, and a full treatment would carry a two-dimensional rate state), from state taxes, from the interaction with wash-sale constraints, and from the step-up at death, which is the ultimate deferral and would extend the horizon to a stochastic mortality date. The certainty-equivalent premium uses a mean-variance approximation that is accurate for small rate variance and less so for large legislated jumps, where the full distribution in Figure 3 is the honest object. Finally, we price the risk but do not solve the hedging problem: there is no traded instrument that pays off on the capital-gains rate, so the practical hedge is realization timing itself, which is what the stopping policy delivers.

8. Conclusion

Deferring a capital gain is usually the right call, and this paper does not dispute that. What it adds is the recognition that the deferral is a position, priceable and occasionally wrong. Holding an unrealized gain is a short position in the future tax rate, worth $g_0(\tau_0 e^{\rho T} - \tau_T)$ at the horizon and carrying a certainty-equivalent haircut of $\frac{\gamma}{2} g_0^2 \text{Var}[\tau_T]$ for the legislative uncertainty. Under mean-reverting rates that haircut is small and the classical defer-everything policy stands. Under a rising risk-neutral drift the realization-timing option acquires a nonempty exercise region, and the rule is to accelerate realizations in low-rate regimes near the horizon, with the trigger given by a single inequality comparing the expected future rate to today's rate grossed up by compounding. For the boutique managing large, low-basis, taxable positions, the number to keep in view is the embedded gain, because both the benefit of deferral and the risk of deferral scale with it, the second one faster than the first.

References

- Arnott, R., Berkin, A., and Ye, J. (2001). Loss Harvesting: What's It Worth to the Taxable Investor? *Journal of Wealth Management*.
- Berkin, A., and Ye, J. (2016). Tax Management, Loss Harvesting, and FIFO Accounting. *Financial Analysts Journal*.
- Chaudhuri, S., Burnham, T., and Lo, A. (2020). An Empirical Evaluation of Tax-Loss-Harvesting Alpha. *Financial Analysts Journal*.
- Constantinides, G. (1983). Capital Market Equilibrium with Personal Tax. *Econometrica*.
- Constantinides, G. (1984). Optimal Stock Trading with Personal Taxes: Implications for Prices and the Abnormal January Returns. *Journal of Financial Economics*.

- Dammon, R., Spatt, C., and Zhang, H. (2001). Optimal Consumption and Investment with Capital Gains Taxes. *Review of Financial Studies*.
- Dammon, R., Spatt, C., and Zhang, H. (2004). Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing. *Journal of Finance*.
- Dybvig, P., and Koo, H. Investment with Taxes. Working paper.
- Feldstein, M., Slemrod, J., and Yitzhaki, S. (1980). The Effects of Taxation on the Selling of Corporate Stock and the Realization of Capital Gains. *Quarterly Journal of Economics*.
- Goldberg, L., Cai, P., and Sosner, N. (2020). Tax-Managed Factor Strategies and the Hurdle from Short-Term Gains. *Financial Analysts Journal / AQR*.
- Haugh, M., Iyengar, G., and Wang, C. Tax-Aware Dynamic Asset Allocation. Working paper.
- Israel, R., Moskowitz, T., Ross, A., and Serban, L. (2020). Tax-Aware Investing: Implications for Factor Portfolios. *AQR*.
- Moehle, N., and Boyd, S. (2020). Tax-Aware Portfolio Construction via Convex Optimization. Working paper, Stanford University.
- Poterba, J. Various work on capital-gains realizations and their elasticity to statutory rates.
- Sialm, C. (2009). Tax Changes and Asset Pricing. *American Economic Review*.
- Sialm, C., and Sosner, N. (2018). Taxes, Shorting, and Active Management. *Financial Analysts Journal*.
- Sosner, N., Gromb, D., Nemtchinov, V., and Pyne, A. (2021). Understanding a Tax-Aware Long-Short Strategy: Separating Alpha from the Tax Benefit. *AQR*.
- Sosner, N., Krasner, S., and Liberman, J. (2019). Loss Harvesting or Gain Deferral? A Surprising Source of Tax Benefits of a Tax-Aware Long-Short Strategy. *Journal of Wealth Management*.
- Stein, D., and Narasimhan, P. (1999). Of Passive and Active Equity Portfolios in the Presence of Taxes. *Journal of Wealth Management*.

Figures




Figure 1: Figure 1: stopping boundary




Figure 2: Figure 2: deferral zero crossing




Figure 3: Figure 3: distribution and haircut