A Statistical Physics Guide to Personal Finance

Vamsi Varanasi

April 2026

Abstract

This writeup arrives at several logical and occasionally well-trodden maxims for how to manage your money by considering wealth as a stochastic jump-diffusion process with dynamics that decouple into three independent regimes: 1) day-to-day (fluctuation-dominated, "quasistatic"), 2) year-to-year (drift dominated), and 3) decade-to-decade (jump-dominated). The dynamics of each regime imply heuristics for decisions made on timescales where the regime applies. Hopefully the quantitative framing makes some ideas click.

1 Setup

Consider your net worth as the following stochastic process [1, 2]:

$$dW = \sigma_W dB_t + \mu dt + dJ_t$$

where

- $\mu = r\bar{W} + (I - C)$ measures how much mean wealth $\bar{W}$ increases per year as a function of $\bar{W}$, income I and costs C
- σ_W is the fluctuation scale of total wealth. This depends on portfolio composition (allocation, leverage), which also sets r, so the two aren't entirely independent.
- $dJ_t = \sum_i Y_i^+ \, dN_t^+ - \sum_j Y_j^- \, dN_t^-$ is a jump term that includes positive (selling your company, inheritance, large windfall) and negative (losing your job, medical issues, market crash) step-wise changes to net worth. $N_t^\pm$ are Poisson processes with rates $\lambda^\pm$, and the jump magnitudes $Y^\pm$ are drawn from some distributions.

Most of these parameters evolve in time and can thus be perturbed by our decisions across various timescales: commonly r, I, C, $\{Y\}$, and $\{N\}$. The latter two evolve in distribution; we usually don't know when jumps happen, but can make decisions to change their likelihood and magnitude (buying more lottery tickets, hedging, etc). Call this set of controls $\theta(t)$.

The game is to choose a trajectory $\theta(t)$ to optimize some objective on W, normally a function of wealth $U(W)$ evaluated at some time t. The right objective is itself a strategic choice and depends on the regime, as we'll see.

1.1 Three regimes

Our SDE has three sources of variation: diffusive fluctuations $\sigma_W \, dB_t$, drift $\mu \, dt$, and jumps dJ_t. Each dominates on a different timescale, with crossovers set by comparing their typical contributions over a window τ.

Fluctuations scale as $\sigma_W \sqrt{\tau}$ and drift as $\mu\tau$, so they cross at

$$\sigma_W \sqrt{\tau} = \mu\tau \implies \tau^* = \left(\frac{\sigma_W}{\mu}\right)^2$$

Below τ^*, fluctuations swamp drift; above it, drift is detectable above the noise.

Assume jumps arrive Poisson with rate λ, so the probability of at least one jump by τ is $1 - e^{-\lambda\tau}$, which becomes order unity at

$$\tau^{**} \sim \frac{1}{\lambda}$$

Whether jumps dominate the dynamics at this scale depends on their size. We assume $\mathbb{E}[Y] \gg \mu\tau^{**}$, i.e. a single jump is much larger than the drift integrated over the inter-jump time (otherwise, jumps are just noise on top of the drift).

These crossovers give:

- $\tau \ll \tau^*$: **fluctuation-dominated.** $\bar{W}$ is approximately constant (quasistatic).
- $\tau^* \ll \tau \ll \tau^{**}$: **drift-dominated.** This is the regime in which most personal finance advice lives.
- $\tau \gtrsim \tau^{**}$: **jump-dominated.** Integrated drift is small compared to the effects of a single jump.

1.2 Toy example

Consider a liquid net worth of $\bar{W} = \$1\text{M}$, all in equities. Annualized equity vol is $\sim$16%, so $\sigma_W \approx 0.16\bar{W} = \$160\text{k}/\sqrt{\text{yr}}$. Take real return $r = 0.07$ and net savings $I - C = \$50\text{k/yr}$, giving $\mu = r\bar{W} + (I - C) = \$120\text{k/yr}$. Then $\tau^* = (160/120)^2 \approx 1.8$ yr.

For jumps, take $\lambda^+ \sim 0.05$/yr with $\mathbb{E}[Y^+] \sim \$3\text{M}$ (a major liquidity event, exit, or inheritance once or twice in a career) and $\lambda^- \sim 0.05$/yr with $\mathbb{E}[Y^-] \sim \$2\text{M}$ (a major medical, legal, or career shock). Total $\lambda = 0.1$/yr, so $\tau^{**} \approx 10$ yr. Both jump magnitudes are well above $\mu\tau^{**} = \$1.2\text{M}$.

Below $\sim$2 years, fluctuations dominate. Between 2 and 10 years, drift dominates. Beyond a decade, jumps dominate, and their sign and size matter more than any decade of saving. We will refer back to this toy example consistently.

1.3 Moving the regime boundaries

Obviously, changing SDE parameters changes regime boundaries. This leads to some interesting dynamics:

- *Early career:* $\bar{W}$ is small and so is $\sigma_W \propto \bar{W}$, so τ^* is much shorter. Drift dominates almost immediately, which is why standard advice works well for young accumulators; most decisions should be evaluated in a drift-centric regime (save money, spend less, optimize taxes).

- *Higher invested wealth at fixed savings rate:* The HENRY's dilemma: $\sigma_W \propto \bar{W}$ while active drift stays fixed, so τ^* grows roughly quadratically in $\bar{W}$. At $\bar{W} = \$5\text{M}$ with the same savings, $\tau^* \approx 5$ yr, compared to $\tau^* \approx 2$ yr at $\bar{W} = \$1\text{M}$. If your income (or rather, your net savings rate) plateaus, fluctuations dominate over longer and longer windows as a function of your wealth, but your wealth might not be high enough to furnish a tolerable Stable Withdrawal Rate [3].

- *Higher leverage:* When you lever a portfolio by L, returns scale by L and so does volatility: $\sigma_W \to L\sigma_W$. The drift also picks up a leverage term, but at the level of order-of-magnitude estimates the dominant effect is on σ_W. Since $\tau^* = (\sigma_W/\mu)^2$, doubling σ_W quadruples τ^*, hence $\tau^* \propto L^2$.

  Concretely: at $L = 2$, the fluctuation regime stretches from $\sim$2 years to $\sim$8 years for the toy example. Leverage doesn't just amplify returns; it pushes you into longer windows of fluctuation-dominated dynamics, which is part of why levered drawdowns are psychologically brutal: drift takes much longer to reassert itself. Note that leverage in formulation encompasses far more things than borrowing money; starting a company is a form of leverage.

- *Higher jump rates:* founders, traders, concentrated positions, but also medical issues, political instability, etc. shrink τ^{**} to a few years; the jump regime arrives early.

Breaking up the optimization problem into regimes also decouples when we care about which control variables. In other words, our policy for the jump parameters $\{Y\}$ and $\{N\}$ doesn't matter in the quasistatic regime, which is somewhat obvious. But less obvious is that our gross savings rate $I - C$ doesn't matter either for purchases smaller than $\sigma_W\sqrt{\tau}$; the entire program is written relative to the noise floor. In the subsequent sections, we'll consider heuristics programs for optimal behavior within each regime, with the upside that we can conduct the programs for each regime in parallel—how we choose whether to buy a coffee at Starbucks on a given morning has little to do with how we maximize our chances of striking gold.

2 The quasistatic regime ($\tau \ll \tau^*$)

In this regime, $\bar{W}$ is approximately constant and the dominant variation is the diffusive term $\sigma_W\sqrt{\tau}$. The optimization problem reduces to a question of detectability: which decisions actually move W enough to be distinguished from noise on the timescale of the decision?

For an expense or gain X on timescale τ, define the signal-to-noise ratio

$$\mathrm{SNR}(X,\tau) = \frac{X}{\sigma_W\sqrt{\tau}}$$

If SNR $\ll 1$, the decision is below the noise floor: whatever you decide, the realized change in W over τ will be dominated by random fluctuations, not your choice. Optimization effort produces nothing observable.

In the toy example, $\sigma_W = \$160\text{k/yr}^{1/2}$. This means:

- Daily ($\tau = 1/365$ yr): $\sigma_W\sqrt{\tau} \approx \$8$k. Anything below ∼\$1k has SNR < 0.1.
- Monthly ($\tau = 1/12$ yr): $\sigma_W\sqrt{\tau} \approx \$46$k. Anything below ∼\$5k has SNR < 0.1.
- Annual ($\tau = 1$ yr): $\sigma_W\sqrt{\tau} = \$160$k. Anything below ∼\$15k has SNR < 0.1.

These are the noise floors. An expense an order of magnitude below them is genuinely indistinguishable from a price tick. Note that they only apply if we're in the critical $\tau < \tau^*$ regime; if τ^* is 6 months, then a \$15k expense could absolutely be material. From these numbers, we can deduce a few maxims:

2.1 Maxims

2.1.1 Don't optimize below the noise floor.

The cognitive cost of deciding whether to spend \$50 vs \$30 on dinner is real; the financial cost is invisible against \$8k of daily fluctuation.

2.1.2 Optimization effort should scale with SNR, not with X.

A \$200 decision and a \$5k decision both feel like "money decisions," but on a daily timescale only the second is real.

2.1.3 The noise floor rises with $\bar{W}$.

Habits formed at low wealth (tracking small expenses, optimizing every transaction) become counterproductive as σ_W grows. At $\bar{W} = \$5$M with the same allocation, the daily noise floor is \$40k, and almost no consumer-level decision clears it.

2.1.4 Recurring decisions live on the longest timescale they recur over, not the shortest.

A single \$1k daily expense is below the daily noise floor; sustained daily \$1k spending is \$365k/yr, which dominates even the annual noise floor and meaningfully shifts μ_{active}. The right X to evaluate is the integrated expense on its natural timescale: a \$50/month subscription is a \$600/yr decision; a daily habit is an annual one. Most lifestyle creep operates this way, hidden as small recurring choices that look like noise individually but are coherent drift in aggregate.

2.1.5 Fluctuation-regime "wins" don't compound.

The converse is also true: a good month in the market or a windfall day doesn't change $\bar{W}$ on any meaningful timescale; it's a draw from $\sigma_W\sqrt{\tau}$. Treating it as a permanent gain (lifestyle creep, "house money" effects) is a category error: you're integrating noise as if it were drift.

The substantive content of this regime is mostly negative: it tells you what *not* to optimize. The active optimization happens in regime 2, which we turn to next.

3 The drift-dominated regime ($\tau^* \ll \tau \ll \tau^{**}$)

In this regime, fluctuations have averaged out and jumps haven't arrived. Net worth evolves approximately deterministically as

$$W(\tau) \approx \bar{W} + \mu\tau, \qquad \mu = r\bar{W} + (I - C)$$

The optimization problem is to tune the components of μ. This is where almost all standard personal finance advice lives: maximize savings, minimize taxes, optimize allocation, avoid lifestyle creep. The advice is correct, but its relevance depends on which component of μ dominates.

3.1 Passive vs active drift.

Decompose $\mu = \mu_{\text{passive}} + \mu_{\text{active}}$ with $\mu_{\text{passive}} = r\bar{W}$ and $\mu_{\text{active}} = I - C$. The two scale differently: passive grows linearly in $\bar{W}$, active is roughly fixed by career and lifestyle (in that if it grows, it grows by our control program, independent of wealth). They cross at a wealth scale

$$W_\times = \frac{I - C}{r}$$

Below $W_\times$, active drift dominates: every dollar of savings or tax optimization moves μ meaningfully. Above $W_\times$, passive drift dominates: a small change in r (allocation, leverage, fee drag) outweighs large changes in savings.

For the toy example with $r = 0.07$ and $I - C = \$50\text{k}$, $W_\times \approx \$700\text{k}$. So at $\bar{W} = \$1\text{M}$ we're already in the passive-dominated subregime, though not by much.

3.2 Maxims

3.2.1 Standard personal finance advice is regime-2 advice and assumes $\bar{W} \ll W_\times$.

"Save more, cut expenses, max your 401(k)" is the right advice when active drift dominates. It's roughly the right advice into the crossover. It becomes increasingly cosmetic above $W_\times$.

3.2.2 Below $W_\times$, optimize active drift.

Savings rate, tax structure (tax-advantaged accounts, deferred comp, QSBS, loss harvesting), and consumption discipline are the highest-leverage controls. A 10% improvement in $I - C$ is worth more than a 1% improvement in r.

3.2.3 Above $W_\times$, optimize passive drift.

Allocation, leverage, fee drag, and tax efficiency on the portfolio dominate. A 1% improvement in r adds $0.01\bar{W}$ to drift annually, which scales without bound. Squeezing more savings out of a fixed income doesn't.

3.2.4 Allocation and leverage set both r and σ_W.

You can't increase μ_{passive} without increasing σ_W, which extends τ^* and pushes you back into the fluctuation regime on longer windows. The right allocation depends on the timescale you actually care about; chasing return without accounting for the timescale extension is how levered drawdowns become psychologically and practically catastrophic.

3.2.5 Correlation between income and portfolio is systematically underweighted.

If I is correlated with the portfolio (tech equity for tech workers, founder equity, financial sector exposure for finance workers), the effective σ_W is larger than naive portfolio vol suggests, and you might actually not be in this regime. Decorrelating the liquid book from the income source is one of the few high-leverage moves available in this regime. This gives us a corollary maxim:

Corollary: Your liquid portfolio's job in regime 2 is partly to hedge your income stream, not just to compound. Hold things that pay off when your job doesn't. The largest decorrelation gains come from sector tilts away from your employer and away from your income source, not from minor allocation tweaks within your concentrated sector [4, 5].

3.2.6 Lifestyle creep is the regime-2 equivalent of integrated noise.

A small permanent increase in C permanently reduces μ_{active}, which compounds forever. The damage isn't the immediate spend; it's the permanent shift in the SDE's drift.

3.3 The drift regime collapses at high enough $\bar{W}$ and/or λ

The drift regime is bounded: its width is $\Delta\tau = \tau^{**} - \tau^*$. As $\bar{W}$ grows at fixed allocation, σ_W and μ both scale linearly, but with different intercepts:

$$\frac{\sigma_W}{\mu} = \frac{\sigma_p \bar{W}}{r\bar{W} + (I - C)} \xrightarrow{\bar{W}\to\infty} \frac{\sigma_p}{r}$$

So τ^* asymptotes from below to a finite limit $(\sigma_p/r)^2$. For equity-like parameters $\sigma_p = 0.16$ and $r = 0.07$, $\tau^* \to 5.2$ yr. The fluctuation regime stretches to about five years and stays there.

Meanwhile $\tau^{**} \sim 1/\lambda$ is set by life events, which we estimate is ∼10 yr for typical careers and shorter for founders, traders, or concentrated positions. The drift window shrinks to $\Delta\tau \sim 1/\lambda - (\sigma_p/r)^2$, and can collapse entirely if jump rates are high.

The drift regime is valid only when $(\sigma_p/r)^2 \ll 1/\lambda$, i.e. when major life events are rare on the scale of portfolio fluctuations. For most accumulators this holds easily and the drift program is the right focus. For high-λ trajectories (active founders, concentrated holders, anyone with frequent large career bets), the drift regime is narrow or nonexistent: they skip directly from fluctuations to jumps, and the optimization shifts accordingly. This is doubly true if you've both already made some money (high $\bar{W}$), which pushes up τ^* concomitantly with high-λ pulling down τ^{**}. The next section takes up the program in the jump-dominated regime.

4 The jump-dominated regime ($\tau \gtrsim \tau^{**}$)

If your jump parameter λ is appreciable, over decades, a few large jumps dominate the trajectory. Drift is small relative to a single $Y^{\pm}$, fluctuations have averaged out, and the integrated W is essentially set by which jumps arrive and when. The optimization problem changes shape: instead of tuning μ, you're shaping the rates $\lambda^{\pm}$ and the distributions of jump magnitudes $Y^{\pm}$.

This dynamic changes our objective function. In regime 2, drift is approximately deterministic and maximizing μ is equivalent to maximizing $\mathbb{E}[W]$. Implicitly, we set our utility function $U(W) = W$, which it is to first order on drift-relevant timescales. But when a single jump can get you far more wealth than drift can, we have to consider the shape of U across orders of magnitude on time scales where jumps are plausible.

Empirically, utility is concave in wealth: new capabilities unlock with orders of magnitude, not at linear increments (e.g. you can't buy a house, then you can; you can't retire, and then you can). Self-reported wellbeing is also concave in income [6, 7]. Both observations point to log utility as a natural caricature: $U(W) \sim \log W$ treats each doubling of wealth as roughly equivalent in capability gain.

The same conclusion follows from a different argument: under multiplicative compounding [8], the long-run growth rate of a single trajectory is $\mathbb{E}[\log W]$, not $\mathbb{E}[W]$, and only the former is what the person living the trajectory actually experiences. Maximizing $\mathbb{E}[W]$ on a compounding process leads to overbetting and eventual ruin even when every individual bet has positive expected value. The two arguments (concave utility and multiplicative compounding) converge on the same objective: maximize $\mathbb{E}[\log W]$.

4.1 Rate-and-magnitude shaping.

Assuming jump magnitudes are multiplicative ($W \to W(1+Y)$ with $Y > -1$), the long-run growth rate decomposes as

$$g \equiv \frac{1}{\tau}\mathbb{E}\left[\log \frac{W(\tau)}{W(0)}\right] = \mu_{\text{eff}} + \lambda^+\mathbb{E}[\log(1+Y^+)] - \lambda^-\mathbb{E}[\log(1-Y^-)]$$

where $\mu_{\text{eff}} = \mu - \frac{1}{2}\sigma_W^2$ is the drift adjusted for diffusive volatility drag, and the jump terms are compound Poisson contributions.

In regime 3, $\mu_{\text{eff}}\tau \ll Y^\pm$ by assumption, so the jump terms dominate. The control problem is to choose $\{\lambda^\pm, Y^\pm\}$ to maximize g, subject to the constraint that modifying any of them costs drift (e.g. paying insurance premiums, sacrificing salary to start a company). The relevant marginal values are

$$\frac{\partial g}{\partial \lambda^+} = \mathbb{E}[\log(1+Y^+)], \qquad -\frac{\partial g}{\partial \lambda^-} = -\mathbb{E}[\log(1-Y^-)] > 0$$

An action that increases λ^+ at drift cost c is worth taking iff $\Delta\lambda^+ \cdot \mathbb{E}[\log(1+Y^+)] > c$. Symmetrically for reducing λ^-.

4.2 Heavy tails dominate

The log in the objective rewards heavy right tails and punishes heavy left tails asymmetrically around 1. Concretely:

- *Heavy-tailed Y^+:* $\mathbb{E}[\log(1+Y^+)]$ can be large even when $\mathbb{E}[Y^+]$ is small or negative, because the right tail contributes to log at any positive value. This is why founding companies, angel investing, and concentrated equity in moonshots are worth pursuing despite negative EV. The "edge" isn't in the mean; it's in the tail.

- *Heavy-tailed Y^-:* $\mathbb{E}[\log(1-Y^-)]$ blows up as $Y^- \to 1$ (full ruin), so any positive probability mass near $Y^- = 1$ contributes infinite penalty per unit rate. Catastrophic medical, legal, or ruin-class events are bounded only by the left tail of Y^-, and the log punishes them disproportionately. This is why true tail-risk insurance is correct even when the premium is much larger than $\mathbb{E}[Y^-]$.

Linear utility ($\mathbb{E}[W]$) is symmetric and would price these bets very differently. Under $\mathbb{E}[\log W]$, both founding-class and insurance-class actions are dramatically more valuable than they look. This is the formal reason expected-value reasoning systematically mismanages regime-3 decisions in both directions.

4.3 Maxims

4.3.1 When the drift regime collapses, optimize jumps, not drift.

The drift regime is meaningful only when $(\sigma_p/r)^2 \ll 1/\lambda$. When that fails, either because $\bar{W}$ has grown enough that τ^* approaches its asymptote, or because λ is high enough that τ^{**} has shrunk, jumps dominate the dynamics regardless of where you are relative to $W_\times$. The relevant question shifts to: which actions modify $\lambda^\pm$ or $Y^\pm$ in ways that move you across orders of magnitude?

4.3.2 Heavy-tailed positive jumps are the most valuable thing in the system.

Founding, angel investing, concentrated equity in moonshots all give large $\mathbb{E}[\log(1+Y^+)]$ at modest rates. Negative EV does not disqualify them; the log rewards the tail.

4.3.3 Heavy-tailed negative jumps are the most dangerous thing in the system.

Catastrophic events with positive mass near $Y^- = 1$ contribute unbounded penalty per unit rate. Pay willingly for protection even at large premiums. Do cardio.

4.3.4 When only jumps matter, pay extra attention to jump dynamics.

Consider that there are many different kinds of jumps. Each will have its own distribution of expected magnitude and frequency, perhaps with some correlations. In particular, pay attention to:

- **Endogenous vs exogenous jumps:** You control rates and magnitudes of some jumps (founding, marriage, geographic moves, concentrated investments) and not others (markets, health, family). Time the endogenous against the exogenous baseline.
- **Reversible vs absorbing jumps:** Some jumps can be unwound at a cost; some cannot. Bias toward reversibility under uncertainty.
- **Clustered jumps:** Marriage, kids, parental care, market regime change, career pivots correlate in time. Plan for joint tails, not marginals.

4.3.5 The right baseline is "W when a jump arrives," not "W at horizon T."

Jumps are Poisson; you don't choose when they hit. This pushes toward higher liquid reserves than pure drift optimization suggests.

4.4 The strategic question

Regime 3 is where personal finance becomes life planning. The controls are slow (you can't change $\lambda^{\pm}$ overnight), the consequences are large, and the optimization is over distributions, not numbers. The right questions are:

- Which order of magnitude am I at, which am I targeting, and what jumps cross between them?
- Which λ's am I paying to modify, and at what cost in drift?
- What's my joint tail with my career, my family, my geography, and the macro?
- Where am I overpaying for tails I don't actually want, or underpaying for ones I do?

Most strategic financial errors are regime-3 errors hidden in regime-2 framing: people optimize savings rates and tax efficiency when the actual question is whether to take a concentrated bet, accept an absorbing risk, or pay for an option. The drift program is real and should run on autopilot once it's tuned. The jump program is where the leverage is.

5 Conclusion

Decisions in life are naturally time bound: buying shoes, taking a job, choosing a career. As we make decisions to accumulate wealth, among other things, we must consider the relevant time scales of the accumulation process to inform our decisions. Our model here details how the time scale cutoffs for each regime vary with total wealth and parameter values, and how the dynamics of different regimes imply different optimal heuristics. Of course, this framework applies to any other stochastic accretion process (wisdom, citations, experiences) where we deem accretion to be a good thing. Feel free to make an appropriate substitution if you find my overtly capitalist framing jarring.

6 Acknowledgements

I thank Claude Opus 4.7 (Anthropic, 2026) for helpful conversations.

References

[1] Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.

[2] Rama Cont and Peter Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2003.

[3] William P. Bengen. Determining withdrawal rates using historical data. *Journal of Financial Planning*, 7(4):171–180, 1994.

[4] Zvi Bodie, Robert C. Merton, and William F. Samuelson. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16(3-4):427–449, 1992.

[5] João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18(2):491–533, 2005.

[6] Daniel Kahneman and Angus Deaton. High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, 107(38):16489–16493, 2010.

[7] Matthew A. Killingsworth, Daniel Kahneman, and Barbara Mellers. Income and emotional well-being: A conflict resolved. *Proceedings of the National Academy of Sciences*, 120(10):e2208661120, 2023.

[8] J. L. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35(4):917–926, 1956.