The Derived Time Stop

Exiting a Swing Trade under Drift Uncertainty and a Deadline

Working Paper · Stochastic Optimal Control

T. Zamrik | 11-JUN-2026

Abstract. A swing trade is a hypothesis with a deadline. The trader enters a setup believing it may carry an edge, the price then either drifts the way the setup promised or it does not, and the setup expires at a known date. We model the drift as one of two constants, unknown to the trader, and ask when the position should be closed and what the usual rules of thumb cost against the answer.

The trader's belief that the edge is real is a filtered diffusion and an explicit function of time and return. Holding on is worth the expected drift under that belief, accumulated until exit, and the problem becomes a finite-horizon optimal stopping problem in the belief alone, with two parameters in information units: the break-even belief θ and the deadline S. Two partial differential equations answer it. A Hamilton–Jacobi–Bellman equation, in the form of a variational inequality solved backward from the deadline, gives the value of holding on, and its free boundary is the exit rule, which says where the trader should get out. A pair of Fokker–Planck equations, one for a real edge and one for no edge, solved forward from entry with that boundary as an absorbing wall, gives the law of the exit time, which says when trades are actually cut. The HJB boundary rises continuously to θ exactly at the deadline, so patience is largest at entry and vanishes at expiry.

The boundary is universal. In log-odds it is $\operatorname{logit}\theta + g(S-s)$ for a single function g of the time left, the same for every break-even belief, because the value of any exit rule is a combination of two expected holding times that do not depend on θ. The function is tabulated once: $g(u)/\sqrt{u}$ stays between -0.63 and -0.58 for time left u between 0.01 and 4. On the return chart the boundary is a stop line that tightens as the deadline nears whenever the average of the two drifts is non-negative, and can first loosen when it is negative.

The Fokker–Planck densities show that the rule cuts 13.8 per cent of real edges and 44.7 per cent of dead trades before the deadline, most of them late, and that a trade cut at time s had a real edge with probability exactly the boundary $b(s)$. A time stop that ignores the price is worth exactly $(p_0 - \theta)$ per unit of time held and can never beat holding to the deadline. The trader's conditional stop, which exits at a fixed day unless the trade has worked, cuts 31.0 per cent of real edges at half the deadline, gives up between 3.3 and 20.6 per cent of the attainable value across the parameters studied, and is worse than simply holding in thirty of the forty-five cases computed. Every number is computed by finite differences for both equations and checked against closed forms, a no-look and a one-look bound, and a million-path simulation.

1. Introduction

Every swing trader carries a time stop, stated or not. The folklore version is short: if the trade has not worked in three days, get out. It sits beside the price stop in trading manuals, it is taught as discipline, and it is almost never derived. The price stop at least has a literature that asks whether it pays: Kaminski and Lo [11] show that a stop-loss rule adds value only when returns carry momentum, and costs value when they are a random walk. The time stop has no such treatment, and it deserves one, because it encodes a claim about information rather than about price: that a setup which has not moved by a certain day has revealed that it carries no edge.

That claim can be right, and it can be badly wrong. A real edge is small against the day-to-day noise of a price, so a trade that has not worked by day three is weak evidence that it never will. A rule that treats it as proof throws away real edges in quantity. This paper measures how many, and what they were worth.

The right way to think about an open swing trade is as a hypothesis under test. At entry the trader holds a belief that the setup carries an edge. Every day of price action is an observation, and the belief is updated by it; the position is worth holding while the updated belief justifies the expected drift, and not otherwise. This is sequential analysis in the sense of Wald and Shiryaev, with the filter of Wonham [1] as its engine: the observation is a Brownian motion with an unknown drift, and the posterior probability of each drift is a diffusion driven by the innovation.

Two features separate the trader's problem from the classical test. The first is that the trader is paid while testing: holding on earns the drift, whatever it is, so the cost of an observation is not a fixed fee but the expected drift under the current belief, which may be positive. The second is the deadline. A setup expires, at an earnings date, at the end of a pattern's window, or when the trader must free the capital, and a finite horizon changes the shape of the answer, as it does in the finite-horizon Wiener test of Gapeev and Peskir [4]. The exit boundary is no longer a level but a curve in time, and the curve is the result.

This paper solves that problem and states the answer in the trader's coordinates. Its contributions are six.

1. **The reduction.** The trade is a stopping problem in the belief alone, with two parameters: the break-even belief θ, below which the expected drift is negative, and the deadline S measured in units of the edge's signal-to-noise (Propositions 2 and 3). No market number has to be assumed to state any result.

2. **The boundary.** The optimal rule exits when the belief falls below a continuous boundary that is below θ before the deadline and rises to θ exactly at it (Theorem 1). Without a deadline there is no stop at all: the boundary falls to zero (Proposition 5).

3. **Universality.** In log-odds the boundary is $\operatorname{logit}\theta + g(S - s)$ for one function g of the time left, the same for every break-even belief (Theorem 2). The whole policy is one curve, tabulated in Table 1.

4. **The stop line.** On the return chart the boundary is an explicit stop line (Corollary 1); it tightens as the deadline nears unless the average of the two drifts is negative, and entry conviction shifts it by exactly the log-odds of the entry belief (Remark 2).

5. **What the rules of thumb cost.** A time stop that ignores the price is worth exactly $(p_0 - \theta)$ per unit of time held (Theorem 3), and the best rule that looks only

once has a closed form (Theorem 4). The trader's conditional time stop is measured against both and against the derived rule over the whole parameter plane.

6. **When the rule exits.** The exit time is a first passage through the HJB boundary, and its density under each hypothesis comes from the forward, Fokker–Planck, equation through the same boundary (Proposition 6). The backward equation says where the trader exits; the forward equation says when.

2. The Setup and the Trader

Notation 1. A filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ carries a standard Brownian motion W and an independent random variable μ taking the value μ_H with probability p_0 and μ_L otherwise. X_t is the trade's log-return since entry and $\mathcal{F}_t^X$ the information it generates. The drift gap is $\Delta = \mu_H - \mu_L > 0$ and the volatility $\sigma > 0$. The break-even belief is $\theta = -\mu_L/\Delta$. Information time is $s = t\,\Delta^2/\sigma^2$, the deadline in the same units is $S = T\Delta^2/\sigma^2$, and the return in units of σ^2/Δ is $x = X\Delta/\sigma^2$. The belief is $\pi_t = \mathbb{P}(\mu = \mu_H \mid \mathcal{F}_t^X)$ and its log-odds $y = \operatorname{logit}\pi = \log(\pi/(1-\pi))$. Time left to the deadline is written $u = S - s$, and Φ is the standard normal distribution function. *Binds.* X, μ, Δ, σ, θ, s, S, x, π, y, u, Φ.

Definition 1 (The return process with an unknown drift)**.** The trade's log-return since entry evolves as

$$dX_t = \mu\,dt + \sigma\,dW_t, \qquad X_0 = 0,$$

where $\mu \in \{\mu_L, \mu_H\}$ is drawn once, before entry, and never observed; $\mu = \mu_H$ with prior probability $p_0 \in (0,1)$. The trader observes X and nothing else. *Binds.* X (the return since entry), μ (the unknown drift, μ_H if the setup carries an edge and μ_L if it does not), σ (the volatility).

The two-point prior is the simplest model in which a trade can be *right or wrong* rather than more or less right, and it is the model a trader actually uses when deciding whether a setup has worked. The two values need not be symmetric: μ_H is the drift of a setup that works as advertised, μ_L the drift of one that does not. Assumption 1 takes $\mu_L < 0$, the case in which a failed setup costs money to hold; with $\mu_L \geq 0$ there is nothing to decide.

Definition 2 (The deadline and the trader's reward)**.** The setup expires at a known time $T < \infty$. An exit rule is an $\{\mathcal{F}_t^X\}$-stopping time $\tau \leq T$, and the position, one unit held from entry until τ, earns

$$R(\tau) = X_\tau.$$

The trader chooses τ to maximise $\mathbb{E}[R(\tau)]$. *Binds.* T (the deadline), τ (the exit time).

A position that is still open at T is closed there; the deadline is part of the setup, not of the trader's preferences. Expected return is the right criterion for a trader who runs many independent setups of the same kind, each small against the book, which is the swing trader's situation. Risk aversion over a single trade is a different problem and is discussed in Remark 4, on the limits of the model.

Assumption 1 (Standing assumptions)**.** Throughout: $\mu_H > 0 > \mu_L$, so that $\theta \in (0,1)$; the volatility σ is constant and known; the position is one unit, entered at time zero and held until τ; there are no transaction costs; and the prior $p_0 \in (0,1)$ is the trader's honest probability that the setup carries an edge. The sign condition is what makes the problem a problem. If $\mu_L \geq 0$ the trade never loses in expectation and the trader should hold to the deadline; if $\mu_H \leq 0$ it should never be entered.

Used by. Every result from Lemma 1 on. Remark 4 records what each assumption buys and what relaxing it would change.

Definition 3 (The belief that the edge is real)**.** The trader's belief at time t is the posterior probability

$$\pi_t = \mathbb{P}(\mu = \mu_H \,|\, \mathcal{F}^X_t), \qquad \pi_0 = p_0,$$

and its log-odds is $y_t = \operatorname{logit} \pi_t$. *Binds.* π (the belief), y (its log-odds).

The expected drift under the belief is $\mu_L + \pi_t \Delta = \Delta(\pi_t - \theta)$, which is positive exactly when $\pi_t > \theta$. That is the only sense in which θ is a break-even point: a trader who could not learn anything more would hold while $\pi > \theta$ and exit otherwise.

Stopping a diffusion whose drift is unknown is a classical problem, and the parts this paper needs are known in pieces. Peskir and Shiryaev [3] give the general machinery of optimal stopping and free-boundary problems, and Gapeev and Peskir [4] solve the Wiener sequential testing problem on a finite horizon, where the deadline turns the classical constant boundaries into curves. In finance the closest problems are about selling. Ekström and Lu [5] sell an asset whose drift is unknown and find the optimal selling time under incomplete information, and Ekström and Vaicenavicius [6] treat liquidation under drift uncertainty with a general prior, showing that the selling boundary is monotone in time. Shiryaev, Xu and Zhou [7] find that under a known drift the optimal rule is to buy and hold or sell at once, depending on the sign of a single number; Dai, Zhang and Zhu [8] trade a regime-switching trend with buy and sell boundaries in the belief; Décamps, Mariotti and Villeneuve [9] study investment timing when the value of the project is unknown; and du Toit and Peskir [10] sell a stock as close as possible to its ultimate maximum.

What is missing is the trader's version of the question and the trader's coordinates. The selling problems value the asset at the moment of sale; a swing trader values the trade by the drift earned while holding it, which makes the reward a running one and the break-even belief a parameter of the problem rather than a consequence of it. The literature states boundaries in the belief; a trader needs them on the return chart, in days. And no one, to our knowledge, has measured the conditional time stop, the rule traders actually use, against the optimal rule. Theorem 2, the universality of the boundary in log-odds, is the structural fact that makes the answer fit in one table; we have not found it stated in this form.

3. The Belief

Lemma 1 (The belief is a diffusion driven by the innovation)**.** *Hypotheses.* Definition 1, Definition 3 and Assumption 1. *Statement.* The process

$$\widehat{W}_t = \frac{1}{\sigma}\Big(X_t - \int_0^t (\mu_L + \pi_r \Delta)\, dr\Big)$$

is an $\{\mathcal{F}_t^X\}$-Brownian motion, and the belief solves

$$d\pi_t = \frac{\Delta}{\sigma}\,\pi_t(1-\pi_t)\, d\widehat{W}_t, \qquad \pi_0 = p_0.$$

In particular π is a bounded martingale in $[0,1]$, and it never reaches 0 or 1 in finite time.

Proof. This is the Wonham filter for a two-state signal that does not switch [1]; Liptser and Shiryaev [2] give the general innovation theorem from which it follows. The innovation $\widehat{W}$ is a continuous $\{\mathcal{F}_t^X\}$-martingale with quadratic variation t, hence a Brownian motion by Lévy's characterisation, because $\mathbb{E}[\mu \mid \mathcal{F}_t^X] = \mu_L + \pi_t\Delta$. The equation for π is the Kushner–Stratonovich equation with a static signal, whose drift vanishes because μ does not move. Boundedness makes π a martingale. The likelihood-ratio derivation that follows obtains the same equation directly from the likelihood ratio, which also shows that $\operatorname{logit}\pi_t$ *is finite for every finite t.* $\square$

The filter, from the likelihood ratio. *Start.* Lemma 1 can be read off the likelihood ratio, which is worth doing because the same calculation gives the explicit form of Proposition 1 and the change of measure behind Theorem 2.

Steps. Under the law $\mathbb{P}_H$ of the return with drift μ_H and the law $\mathbb{P}_L$ with drift μ_L, Girsanov's theorem gives the density of one against the other on $\mathcal{F}_t^X$:

$$L_t = \frac{d\mathbb{P}_H}{d\mathbb{P}_L}\Big|_{\mathcal{F}_t^X} = \exp\Big(\frac{\Delta}{\sigma^2}X_t - \frac{\mu_H^2-\mu_L^2}{2\sigma^2}\,t\Big). \tag{3.1}$$

Bayes' formula then gives $\pi_t = p_0L_t/(p_0L_t + 1 - p_0)$, so the log-odds is

$$y_t = \operatorname{logit} p_0 + \Lambda_t, \qquad \Lambda_t = \log L_t = \frac{\Delta}{\sigma^2}X_t - \frac{\Delta(\mu_H+\mu_L)}{2\sigma^2}\,t. \tag{3.2}$$

Itô's formula applied to $\pi = \operatorname{expit}(y)$, with $d\Lambda_t = \frac{\Delta}{\sigma^2}\,dX_t - \frac{\Delta(\mu_H+\mu_L)}{2\sigma^2}\,dt$ and $d\langle\Lambda\rangle_t = \Delta^2\sigma^{-2}\,dt$, gives

$$d\pi_t = \pi_t(1-\pi_t)\,dy_t + \tfrac{1}{2}\pi_t(1-\pi_t)(1-2\pi_t)\frac{\Delta^2}{\sigma^2}\,dt = \frac{\Delta}{\sigma}\pi_t(1-\pi_t)\,d\widehat{W}_t, \tag{3.3}$$

after substituting $dX_t = (\mu_L + \pi_t\Delta)\,dt + \sigma\,d\widehat{W}_t$ and collecting the drift terms, which cancel.

End. The belief is a smooth function of (t, X_t) alone, and the log-likelihood ratio Λ is a Brownian motion with drift $+\frac{1}{2}\Delta^2/\sigma^2$ under $\mathbb{P}_H$ and $-\frac{1}{2}\Delta^2/\sigma^2$ under $\mathbb{P}_L$. Neither law involves θ or p_0; that is the fact Theorem 2 turns into a theorem.

Proposition 1 (The belief is an explicit function of time and return)**.** *Hypotheses.* Lemma 1. *Statement.* In information units,

$$y_s = \operatorname{logit}\pi_s = \operatorname{logit}p_0 + x_s - \left(\tfrac{1}{2} - \theta\right)s,$$

so the belief at any moment is read from the return since entry and the time since entry, with no memory of the path between.

Proof. Substitute $x = X\Delta/\sigma^2$ and $s = t\Delta^2/\sigma^2$ into the log-likelihood ratio of the derivation after Lemma 1. The drift term becomes $\frac{\mu_H+\mu_L}{2\Delta}\,s$, and since $\mu_L = -\theta\Delta$ and $\mu_H = (1-\theta)\Delta$ its coefficient is $\frac{1-2\theta}{2} = \frac{1}{2} - \theta$. The statement has a trading reading that is easy to miss. The belief rises with the return, one-for-one in log-odds, and drifts down at the rate $\frac{1}{2} - \theta$ when $\theta < \frac{1}{2}$: a trade that merely stands still is slowly losing the argument, because a setup with an edge would have been expected to move. When $\theta > \frac{1}{2}$ the average drift is negative and standing still is mild evidence for the edge. □

4. The Stopping Problem

Definition 4 (The value of holding on)**.** The expected drift per unit time under the belief is

$$m(\pi) = \mu_L + \pi\Delta = \Delta(\pi - \theta),$$

and the value of holding an open position at time t with belief π is

$$V(t,\pi) = \sup_{t\le\tau\le T} \mathbb{E}\Big[\int_t^\tau m(\pi_r)\,dr \;\Big|\; \pi_t = \pi\Big],$$

the supremum over exit times of the observation filtration. *Binds.* V (the value of holding on), m (the expected drift under the belief).

Exiting at once is always available, so $V \ge 0$; the trader is in the market exactly where $V > 0$. Proposition 2 shows that V is also the value of the original problem, so nothing is lost by working with the belief.

Proposition 2 (The problem is a stopping problem in the belief alone)**.** *Hypotheses.* Definitions 2 and 4, Lemma 1. *Statement.* For every exit time $\tau \le T$,

$$\mathbb{E}[X_\tau] = \mathbb{E}\int_0^\tau m(\pi_t)\,dt,$$

so the trader's problem of Definition 2 has value $V(0,p_0)$, and an exit time is optimal for one problem exactly when it is optimal for the other.

Proof. By Lemma 1, $X_t = \int_0^t(\mu_L + \pi_r\Delta)\,dr + \sigma\hat{W}_t$ with $\hat{W}$ an $\{\mathcal{F}_t^X\}$-Brownian motion. For a bounded stopping time optional sampling gives $\mathbb{E}[\hat{W}_\tau] = 0$, which is the identity. The right-hand side depends on the path only through π, a time-homogeneous Markov diffusion, *so the problem is an optimal stopping problem for π with running reward m in the sense of [3]. The reduction replaces a quantity the trader cannot see, the true drift, with one the trader can compute at every moment. It also explains why the problem is not trivial. Without learning, the reward rate $m(\pi_t)$ would be frozen at $m(p_0)$ and the answer would be all-or-nothing; learning makes the rate move, and the option to leave when it turns negative is what has value.* □

Proposition 3 (The problem in information units: the break-even belief and the deadline)**.** *Hypotheses.* Proposition 2. *Statement.* With $s = t\Delta^2/\sigma^2$ and $S = T\Delta^2/\sigma^2$,

$$V(t,\pi) = \frac{\sigma^2}{\Delta}\, v(s,\pi), \qquad v(s,\pi) = \sup_{s \le \tau \le S} \mathbb{E}\Big[\int_s^\tau (\pi_r - \theta)\, dr\Big], \qquad d\pi_s = \pi_s(1-\pi_s)\, dB_s,$$

for a standard Brownian motion B in information time. The optimal exit rule depends on the market only through θ and S, and on the trader only through p_0. In log-odds the belief moves as $dy_s = dB_s + (\pi_s - \frac{1}{2})\, ds$.

Proof. Time-change Lemma 1 by $s = t\Delta^2/\sigma^2$; the diffusion coefficient $\frac{\Delta}{\sigma}\pi(1-\pi)$ becomes $\pi(1-\pi)$, and $m(\pi)\, dt = \Delta(\pi-\theta)\frac{\sigma^2}{\Delta^2}\, ds = \frac{\sigma^2}{\Delta}(\pi-\theta)\, ds$. The log-odds equation is Itô's formula for logit. *The unit of time is the time it takes the edge to show itself above the noise: one information unit is $(\sigma/\Delta)^2$ calendar units, so a setup whose edge is a third of its daily volatility needs nine trading days per unit. The unit of return, σ^2/Δ, is the move that shifts the log-odds by one. Every figure and table in this paper is in these units, and nothing in them depends on a choice of market parameters.* □

Remark 1 (The myopic rule, and why it exits too early)*.* A trader who ignores learning holds while the expected drift is positive, $\pi > \theta$, and exits otherwise. Proposition 2 shows that this rule is wrong at every moment before the deadline. At $\pi = \theta$ the reward rate is zero, but holding for a short time δ and then leaving if the belief has fallen is worth

$$\mathbb{E}\big[(\pi_{s+\delta} - \theta)^+\big](S - s - \delta) > 0,$$

because the belief is a martingale with positive variance. So the break-even belief lies strictly inside the region where holding is worth something, and the optimal rule waits below it.

The gap between the two rules is the value of the option to learn. It is largest at entry, when there is the most time left to act on what is learned, and it vanishes at the deadline, where the two rules coincide. Theorem 1 turns this into a boundary; Proposition 4 measures its size.

5. The Exit Boundary

Theorem 1 (A continuous exit boundary that rises to the break-even belief at the deadline)**.** *Hypotheses.* Proposition 3, with $\theta \in (0,1)$ and $S < \infty$. *Statement.* There is a continuous non-decreasing function $b : [0, S) \to (0, \theta)$ with $\lim_{s \uparrow S} b(s) = \theta$ such that

$$\tau^* = \inf\{r \geq s : \ \pi_r \leq b(r)\} \wedge S$$

is an optimal exit time from every state (s, π). The continuation region is $\{(s, \pi) : \pi > b(s)\}$; there v is $C^{1,2}$ and solves $\partial_s v + \frac{1}{2}\pi^2(1-\pi)^2 \partial_{\pi\pi} v + \pi - \theta = 0$, it vanishes on and below the boundary, and it meets zero smoothly: $\partial_\pi v(s, b(s)+) = 0$.

Proof. We follow the standard route for finite-horizon problems with a running reward [3], [4] and give the steps that are specific to this reward. Monotone in the belief. The flow of $d\pi = \pi(1-\pi)\, dB$ *preserves order, so a higher starting belief gives a pathwise higher belief for every exit rule, and the reward* $\pi - \theta$ *is increasing. Hence* $v(s, \cdot)$ *is non-decreasing and the set where it vanishes is an interval* $[0, b(s)]$*. Monotone in time. The problem is time-homogeneous, so* $v(s, \pi) = w(S - s, \pi)$ *with* w *non-decreasing in the time left, because a longer horizon offers a larger set of exit times. The stopping set therefore shrinks as the time left grows, which is* b *non-decreasing in* s*. Below the break-even belief. Remark 1 shows* $v(s, \theta) > 0$ *for* $s < S$*; continuity of* v *gives* $b(s) < \theta$*. Above zero. Proposition 5 exhibits, for each* $s < S$*, beliefs so low that no rule recovers the expected loss of waiting, so* $b(s) > 0$*. At the deadline. For* $\pi < \theta$ *and* $S - s = \varepsilon$ *small, waiting costs at least* $(\theta - \pi)\varepsilon$ *up to terms of order* $\varepsilon^{3/2}$*, while the most that learning can recover is bounded by* $\mathbb{E}\int_s^S (\pi_r - \theta)^+ \, dr = O(\varepsilon^{3/2})$ *because* $(\pi_r - \theta)^+$ *is zero until the belief has moved by* $\theta - \pi$*. So* $v(s, \pi) = 0$ *near the deadline, and* $b(S-) \geq \pi$ *for every* $\pi < \theta$*. Continuity of* b*, the regularity of* v *in the continuation region and smooth fit follow from the local regularity of the generator away from* $\pi \in \{0, 1\}$ *exactly as in [4].* □

Theorem 2 (One exit curve for every break-even belief)**.** *Hypotheses.* Theorem 1 and Proposition 1. *Statement.* There is one continuous decreasing function $g : (0, \infty) \to (-\infty, 0)$ with $g(0+) = 0$, which does not depend on θ, p_0 or S, such that

$$\operatorname{logit} b(s) = \operatorname{logit} \theta + g(S - s).$$

Equivalently, the trader exits at the first moment the log-odds that the edge is real exceed the log-odds of break-even by less than g of the time left. Moreover $v(s, \pi) = (1 - \pi)\, \theta\, F(\operatorname{logit} \pi - \operatorname{logit} \theta,\ S - s)$ for one function $F \geq 0$, so the value function is universal up to the same factor.

Proof. The holding-time derivation after this theorem shows that, from belief π *at time* s*, any exit rule earns*

$$(1 - \pi)\, \theta \left[e^z\, \mathbb{E}_H[\tau - s] - \mathbb{E}_L[\tau - s] \right], \qquad z = \operatorname{logit} \pi - \operatorname{logit} \theta,$$

where $\mathbb{E}_H$ *and* $\mathbb{E}_L$ *are expectations under the two drifts. Every exit rule is a functional of the future log-odds path, and by Proposition 1 that path is* $z + \operatorname{logit} \theta$ *plus the log-likelihood ratio, whose law under* $\mathbb{P}_H$ *and* $\mathbb{P}_L$ *is a Brownian motion with drift* $\pm\frac{1}{2}$ *that involves neither* θ *nor* p_0*. So after the shift by* $\operatorname{logit} \theta$ *the supremum of the bracket over exit rules is a function* $F(z, S - s)$ *alone, and* v *vanishes exactly where* F *does. Theorem 1 says* *this set is $\{\pi \le b(s)\}$, so it is $\{z \le g(S-s)\}$ for a function g of the time left only. The properties of g are those of b in Theorem 1: decreasing in the time left because b rises with s, negative because $b < \theta$, and tending to zero at the deadline because $b(S-) = \theta$. The theorem is why the whole policy fits in one table. A trader needs g once, as a function of the time left, and every setup, whatever its break-even point, reads its exit boundary from the same curve.* □

Why the break-even belief only shifts the curve: two expected holding times. *Start.* Theorem 2 rests on one identity: the value of any exit rule splits into two expected holding times, one for each drift, that do not depend on where the break-even point is.

Steps. In information units the return has drift μ/Δ, equal to $1-\theta$ when the edge is real and $-\theta$ when it is not, so by Proposition 2 an exit rule started at time s with belief π earns

$$\mathbb{E}\int_s^\tau (\pi_r - \theta)\,dr = \mathbb{E}\Big[\frac{\mu}{\Delta}\,(\tau - s)\Big] = \pi(1-\theta)\,\mathbb{E}_H[\tau - s] - (1-\pi)\,\theta\,\mathbb{E}_L[\tau - s]. \tag{5.1}$$

Factoring out $(1-\pi)\theta$ and writing the odds ratio as an exponential,

$$\frac{\pi(1-\theta)}{(1-\pi)\theta} = \exp\big(\operatorname{logit}\pi - \operatorname{logit}\theta\big) = e^{z}, \tag{5.2}$$

gives the bracket of Theorem 2. The first expectation rewards staying in when the edge is real; the second charges for staying in when it is not.

End. The break-even belief has left the dynamics entirely. It appears only through z, the distance in log-odds between the current belief and break-even, and through the positive factor $(1-\pi)\theta$, which scales the value without moving the boundary.

The free-boundary problem and smooth fit. *Start.* Theorem 1 characterises the value function as the solution of a free-boundary problem, and that problem is what Algorithm 1 solves.

Steps. On the continuation region the value satisfies the backward equation of the belief with the running reward, and on the stopping region it is zero:

$$\partial_s v + \tfrac{1}{2}\pi^2(1-\pi)^2\,\partial_{\pi\pi} v + (\pi - \theta) = 0 \quad \text{for } \pi > b(s), \qquad v = 0 \quad \text{for } \pi \le b(s), \tag{5.3}$$

with the terminal condition $v(S,\pi) = 0$ and smooth fit $\partial_\pi v(s, b(s)+) = 0$. The two regimes combine into a single variational inequality, which is the form a numerical scheme can use without knowing b in advance [3]:

$$\min\Big\{ -\partial_s v - \tfrac{1}{2}\pi^2(1-\pi)^2\,\partial_{\pi\pi} v - (\pi - \theta),\ v \Big\} = 0. \tag{5.4}$$

In log-odds the degenerate coefficient $\pi^2(1-\pi)^2$ disappears. With $\pi = \operatorname{expit}(y)$ the operator becomes $\frac{1}{2}\partial_{yy} + (\pi - \frac{1}{2})\partial_y$, a uniformly parabolic equation on the whole line with a bounded drift.

End. The log-odds form is the one to solve. It has constant diffusion, it resolves beliefs near zero as finely as beliefs near one half, and the boundary it produces is directly the one Theorem 2 is stated in.

Proposition 4 (Patience shrinks as the deadline nears and grows with the deadline)**.** *Hypotheses.* Theorems 1 and 2. *Statement.* The patience $\theta - b(s)$, the distance by which the exit boundary sits below break-even, depends only on the time left and on θ. It is non-increasing in s, vanishes at the deadline, and at entry it is non-decreasing in the deadline S.

Proof. By Theorem 2, $b(s) = \operatorname{expit}\big(\operatorname{logit}\theta + g(S-s)\big)$*, a function of* θ *and the time left alone. It is increasing in* s *because* g *is decreasing, equal to* θ *at the deadline because* $g(0+) = 0$*, and at* $s = 0$ *it is decreasing in* S *for the same reason. The measured sizes are worth having in view. With* $\theta = 0.3$ *the patience at entry is* 0.062, 0.085, 0.114, 0.149 *and* 0.188 *for deadlines* $S = 0.25$, 0.5, 1, 2 *and* 4*. Because the boundary depends only on the time left, the boundary half way through a deadline of* 2 *must equal the boundary at entry for a deadline of* 1*; the solver returns* 0.18628 *for both, and the same identity holds to five digits between each adjacent pair of deadlines.* □

Proposition 5 (Without a deadline there is no stop: the boundary falls to zero)**.** *Hypotheses.* Proposition 3. *Statement.* For every s and every belief $\pi > 0$, $v(s,\pi) \to \infty$ as $S \to \infty$. Hence $b(s) \to 0$: without a deadline no belief is low enough to exit. For finite S the boundary obeys the explicit bound

$$\frac{b(s)}{1-b(s)} \;\leq\; \frac{2}{1-e^{-1}} \cdot \frac{\theta}{1-\theta} \cdot \frac{1}{S-s}.$$

Proof. Take the rule that holds until the log-likelihood ratio has fallen by one unit, or to the deadline. Under $\mathbb{P}_H$ *the ratio is a Brownian motion with drift* $+\frac{1}{2}$*, so it never falls by one with probability* $1 - e^{-1}$*, and on that event the rule holds for the whole of* $S - s$*. Under* $\mathbb{P}_L$ *the drift is* $-\frac{1}{2}$ *and the expected time to fall by one is* 2*. By the holding-time derivation after Theorem 2 the rule earns at least*

$$(1-\pi)\theta\Big[e^z(1-e^{-1})(S-s) - 2\Big], \qquad z = \operatorname{logit}\pi - \operatorname{logit}\theta,$$

which is positive, so π *lies in the continuation region, as soon as* $e^z(1-e^{-1})(S-s) > 2$*; this is the bound. Letting* $S \to \infty$ *at fixed* π *sends the earnings to infinity. The literature on selling under drift uncertainty discounts or caps the horizon [6], [7], and this proposition says why something must. A real edge pays for as long as the trade is held, so with no end date the option to keep learning is worth more than any expected loss, and the trader who never cuts a loser is, in this model, right. The deadline is the only reason a stop exists.* □

Corollary 1 (The stop line in return and days, and when it tightens)**.** Under Theorem 2 and Proposition 1 the trader exits at the first moment the return since entry falls to the stop line

$$x_b(s) = \operatorname{logit}\theta - \operatorname{logit} p_0 + g(S-s) + \left(\tfrac{1}{2} - \theta\right) s,$$

whose slope is $x_b'(s) = -g'(S-s) + \frac{1}{2} - \theta$. Since g is decreasing, the first term is positive. If $\theta \le \frac{1}{2}$ the line tightens throughout the life of the trade. If $\theta > \frac{1}{2}$ it loosens exactly while $|g'(S-s)| < \theta - \frac{1}{2}$, and tightens after that.

Proof. Substitute Theorem 2 into Proposition 1 and set $y = \operatorname{logit} b(s)$. *The slope is the derivative of the right-hand side. In the trader's language, the line has two parts. The log-odds part always tightens, because patience runs out. The drift part moves with the average of the two candidate drifts, which is* $\frac{1}{2} - \theta$ *in these units: when that average is positive a trade that stands still is losing ground, and the stop rises to meet it; when it is negative, standing still is mild evidence for the edge and the stop can fall back. With* $g(u) \approx -0.61\sqrt{u}$ *the loosening lasts while the time left exceeds roughly* $\left(0.305/(\theta - \frac{1}{2})\right)^2$. □

Example 1 (A long deadline: the boundary sinks toward zero, and how fast)**.** *Parameters.* $\theta = 0.3$, deadlines $S = 1$, 4, 16, 64 and 256, solved by Algorithm 1 on a log-odds grid wide enough to hold the boundary. *Computation.* The boundary at entry is $b(0) = 0.1864$, 0.1125, 0.0455, 0.0131 and 0.0034. The bound of Proposition 5 gives 0.5755, 0.2532, 0.0781, 0.0207 and 0.0053, and holds at every deadline.

Worked numbers. Between $S = 16$ and $S = 256$ the logarithm of $b(0)$ falls against the logarithm of S with slope -0.936, and at the longest deadline the measured boundary is 64 per cent of the bound. Both say the same thing: the exit belief falls in proportion to $1/S$.

Shows. Sixteen information units is a long runway: sixteen times the time it takes the edge to show itself above the noise. At that horizon the trader should still be holding when the belief that the edge is real has fallen to four and a half per cent. That is not recklessness; it is the arithmetic of a long runway, in which a real edge that survives earns for so long that it pays for many dead trades held to their proof.

Remark 2 (Conviction at entry buys room: the stop line shifts by the log-odds of the entry belief). By Corollary 1 the entry belief enters the stop line only through the constant $-\operatorname{logit} p_0$. Raising conviction moves the whole line down by exactly the change in log-odds, and changes nothing else about its shape. With $\theta = 0.3$ and $S = 1$ the line starts at $x_b(0) = -0.628$, -1.070, -1.475, -1.880 and -2.322 for $p_0 = 0.3$, 0.4, 0.5, 0.6 and 0.7, and each step is the step in $\operatorname{logit} p_0$.

A setup is worth entering only if p_0 exceeds the entry boundary, $b(0) = 0.186$ here. The first entry in that list is the instructive one. At $p_0 = \theta = 0.3$ the expected drift is exactly zero, holding to the deadline is worth nothing, and yet the derived rule is worth 0.044. The whole of that value is the option to learn and to leave.

6. Computing the Boundary

Algorithm 1 — The projected implicit scheme for the obstacle problem

Inputs. The break-even belief θ, the deadline S, a uniform log-odds grid $y_1 < \cdots < y_N$ on $[-L, L]$ with spacing δy, and M time steps of size $\delta s = S/M$.

Steps.

```
pi <- expit(y);  v <- 0                        # value at the deadline
A  <- tridiag: (1/2) d2/dy2 + (pi-1/2) d/dy    # central differences
for n = M-1 down to 0:
    s   <- n * ds
    rhs <- v + ds * (pi - theta)               # the running reward
    v(-L) <- 0                                 # certain: no edge
    v(+L) <- (pi(+L) - theta) * (S - s)        # certain: an edge
    v   <- solve((I - ds*A) v = rhs)           # one implicit step
    v   <- max(v, 0)                           # exiting is allowed
    i   <- first node with v > 0
    y_b(s) <- y_i - sqrt(v_i)*dy
               / (sqrt(v_{i+1}) - sqrt(v_i))   # smooth fit
return v and y_b(s);  b(s) = expit(y_b(s))
```

Cost. One tridiagonal solve per step, so $O(NM)$ operations and $O(N)$ memory. At $N = 9601$ and $M = 20\,000$ a solve takes a few seconds on one core. The central difference is monotone because $|\pi - \frac{1}{2}|\,\delta y < 1$ on every grid used, and a monotone, stable, consistent scheme converges to the viscosity solution of the variational inequality in the free-boundary derivation [13]; the projected implicit step is the standard treatment of the obstacle for American-type problems [12]. The boundary is located to sub-grid accuracy from the square root of the value, which is linear across the boundary because of smooth fit.

Algorithm 2 — Simulating the rule and the fixed time stops

Inputs. θ, S, p_0, the boundary y_b from Algorithm 1, a number of paths n, a time step δs, and look days h for the fixed stops.

Steps.

1. Draw the truth for each path: the edge is real with probability p_0, and the drift is then $1-\theta$, otherwise $-\theta$.
2. Advance every path's return by $x \leftarrow x + \text{drift} \cdot \delta s + \sqrt{\delta s}\,\xi$ with ξ standard normal, one common noise stream for every rule.
3. Compute the log-odds from Proposition 1, $y = \operatorname{logit} p_0 + x - (\frac{1}{2} - \theta)s$, and close every open path whose y has fallen below $y_b(s) + 0.5826\sqrt{\delta s}$; its payoff is its return at that moment. The shift is the continuity correction for checking a continuous boundary only on a grid, and Proposition 6's forward equation confirms it (see the exit-density check).
4. At each look day h, apply the trader's rule, which continues only if $x_h > 0$, and the best single look of Theorem 4, which continues only if $y_h > \operatorname{logit}\theta$. A path that continues is held to the deadline.

5. `Estimate each rule's value as the mean payoff. Because every rule is run on the same paths as holding to the deadline, whose value` $(p_0 - \theta)S$ `is known exactly, report the paired difference from holding plus that constant: a control variate that removes most of the sampling noise.`

`Cost.` $O(n\, S/\delta s)$. `A million paths at` $\delta s = 10^{-3}$ `over` $S = 1$ `take under a minute in vectorised NumPy. Without the continuity correction the exits come slightly late and slightly too few: half a percentage point of real edges at` $\delta s = 5 \times 10^{-4}$.

7. Verification

The boundary under grid refinement. *Method.* Algorithm 1 is run at four log-odds resolutions with the time step fixed at $\delta s = 10^{-4}$, and at four time steps with the resolution fixed at $N = 4801$, all at the base point $\theta = 0.3$, $S = 1$, $p_0 = 0.5$. Theorem 1's claims are tested at every run: the value at entry, the boundary at entry, the boundary at the deadline, and the largest downward step in the computed boundary, which Theorem 1 says must vanish.

Parameters. Log-odds window $[-18, 18]$; $N = 1201$, 2401, 4801 and 9601 nodes; time steps 4×10^{-4}, 2×10^{-4}, 10^{-4} and 5×10^{-5}.

Result.

grid	δy	$v(0, p_0)$	$b(0)$	**largest downward step**
$N = 1201$	0.0300	0.2076714	0.18582	4.1×10^{-3}
$N = 2401$	0.0150	0.2076710	0.18596	2.2×10^{-3}
$N = 4801$	0.0075	0.2076709	0.18628	9.6×10^{-4}
$N = 9601$	0.0038	0.2076708	0.18624	5.5×10^{-4}

Refining the time step at $N = 4801$ moves the value by less than 5×10^{-8} and the boundary at entry between 0.18586 and 0.18713. At every resolution the boundary at the deadline is $\theta = 0.3$.

Error. The value at entry is converged to seven digits. The boundary is located only to about $\pm 5 \times 10^{-4}$, and it is the boundary, not the value, that carries the grid: it moves between nodes as time passes, and the square-root location of Algorithm 1 is exact only for an exactly quadratic contact. The downward steps are that effect and nothing else. They halve each time the grid is halved, from 4.1×10^{-3} to 5.5×10^{-4}, which is the signature of a discretisation error of order δy rather than of a real turn in the boundary, and Theorem 1's monotonicity holds to within them.

The boundary at entry as the deadline grows. *Method.* Algorithm 1 at $\theta = 0.3$ for deadlines $S = 1$, 4, 16, 64 and 256, on a log-odds window $[-24, 24]$ so that the boundary stays inside the grid as it sinks, compared with the explicit bound of Proposition 5 and fitted in logarithms between $S = 16$ and $S = 256$.

Parameters. $N = 4801$; $\delta s = \min(10^{-4}S,\ 2 \times 10^{-3})$.

Result. The boundary at entry is 0.1864, 0.1125, 0.0455, 0.0131 and 0.0034, against the bound 0.5755, 0.2532, 0.0781, 0.0207 and 0.0053. The fitted slope of $\log b(0)$ against $\log S$ is -0.936.

Error. The bound holds at all five deadlines, and the ratio of the measured boundary to the bound settles near 0.64, which is what Proposition 5 predicts if both fall like $1/S$. A slope slightly shallower than -1 is expected over a finite range because the bound is not sharp; the claim verified is that the boundary falls to zero as the deadline grows, at the rate the proof gives.

Monte Carlo against the PDE value at entry. *Method.* Algorithm 2 follows the boundary of Algorithm 1 on a million paths and estimates the value of the derived rule, which Proposition 2 says must equal the PDE value $v(0, p_0)$. A second test checks Proposition 3 itself: the same exit rule is simulated in two *dimensional* parameterisations with different volatility and drift gap but the same θ and S, and their expected returns, divided by σ^2/Δ, must agree.

Parameters. $\theta = 0.3$, $S = 1$, $p_0 = 0.5$; 10^6 paths at time steps 2×10^{-3}, 10^{-3} and 5×10^{-4}. For the scaling test, $(\sigma, \Delta, T) = (1, 1, 1)$ and $(0.5, 2, 0.0625)$, which both give $S = 1$, on 4×10^5 paths.

Result. The PDE value is 0.207671. The paired simulation estimates are 0.20792 ± 0.00032, 0.20752 ± 0.00032 and 0.20785 ± 0.00032 at the three step sizes; the unpaired means are 0.2074, 0.2080 and 0.2067 with standard error 0.0011. In the scaling test both parameterisations return 0.20693 ± 0.00168 in the scaled unit.

Error. All three paired estimates lie within 0.8 standard errors of the PDE value, and the step size shows no trend at this precision, and the boundary is checked on the grid with the continuity correction of Algorithm 2, so no step-size trend is expected and none appears. The control variate cuts the standard error by a factor of 3.3. The two dimensional runs agree to every printed digit because they share one noise stream, which shows that the change of units of Proposition 3 is an identity path by path, not merely in law; their common value is 0.4 standard errors from the PDE.

The collapse of every boundary onto one curve, and its square-root shape. *Method.* Algorithm 1 is solved for five break-even beliefs spanning almost the whole range, and each boundary is shifted by $-\operatorname{logit}\theta$. Theorem 2 says the five shifted curves are one curve. The spread between them is measured at three resolutions, and the common curve g is then tabulated from a solve with a long deadline.

Parameters. $\theta = 0.1$, 0.3, 0.5, 0.7 and 0.9; $S = 4$; $N = 2401$, 4801 and 9601; the spread taken over 79 values of the time left between 0.05 and 3.95. The table is from $\theta = 0.5$, $S = 16$, $N = 9601$ on $[-24, 24]$.

Result. The largest spread between the five shifted curves is 7.3×10^{-3}, 4.3×10^{-3} and 3.0×10^{-3} at the three resolutions. Over the time left from 0.01 to 4, $g(u)/\sqrt{u}$ lies between -0.627 and -0.579; it is -0.588 at $u = 8$ and -0.553 at $u = 16$.

Error. The spread is below the grid spacing at every resolution and falls as the grid is refined, so the five boundaries are one curve up to discretisation. Nothing in the solver knows the theorem: each break-even belief is solved from its own equation, with its own reward. The square-root shape is a measurement, not a result. It holds to within six per cent of -0.61 over the range where swing trades live, and bends below it for long deadlines, where Proposition 5 says the boundary must fall like $1/S$ and g must therefore grow like $\log u$ rather than $\sqrt{u}$.

Script. `fig_universal.py` — Algorithm 1 for five break-even beliefs at $S = 4$, $N = 9601$.

Shows. Left, the five exit boundaries in log-odds against the time left: five different curves, one for each break-even point. Right, the same five curves shifted by $\operatorname{logit}\theta$, drawn with decreasing line widths so that the overlap can be seen: they are a single curve, g, which starts at zero at the deadline and falls to -1.22 with four units left. The dashed red curve is $-0.61\sqrt{u}$, drawn for scale.

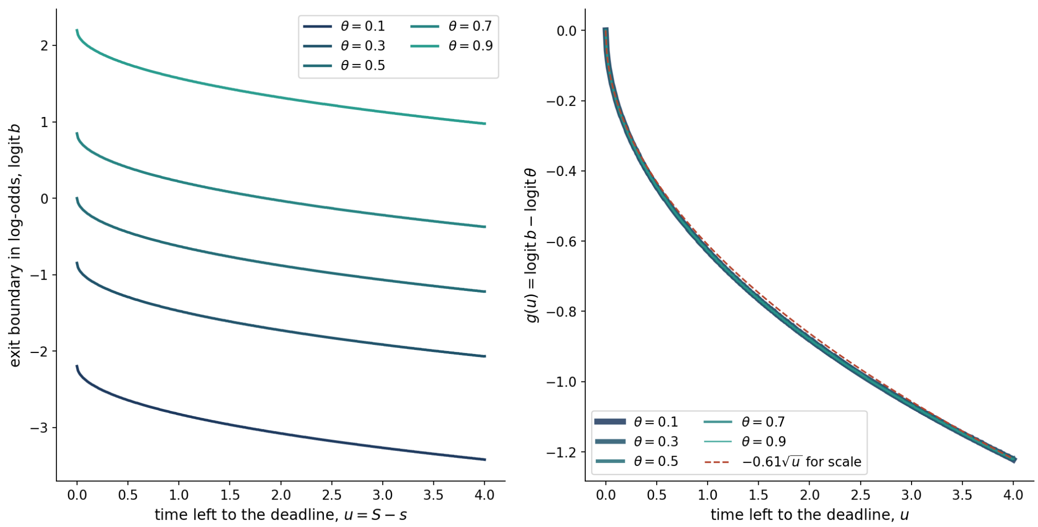

Figure 1: Six break-even beliefs, one exit curve: every exit boundary, shifted by the log-odds of its own break-even belief, falls onto the single function g of the time left

Source. `compute_numbers.py`, section 11 — Algorithm 1 at $\theta = 0.5$, $S = 16$, $N = 9601$ on the log-odds window $[-24, 24]$; the boundary read at ten values of the time left.

Columns. the time left to the deadline u, in information units; the universal curve $g(u)$; its ratio to $\sqrt{u}$; and the exit belief $b = \operatorname{expit}(\operatorname{logit}\theta + g(u))$ it implies for two break-even beliefs, $\theta = 0.3$ and $\theta = 0.7$.

time left u	$g(u)$	$g(u)/\sqrt{u}$	**exit belief,** $\theta = 0.3$	**exit belief,** $\theta = 0.7$
0.01	−0.0579	−0.579	0.2880	0.6877
0.05	−0.1360	−0.608	0.2722	0.6707
0.1	−0.1943	−0.614	0.2608	0.6577
0.25	−0.3121	−0.624	0.2388	0.6307
0.5	−0.4431	−0.627	0.2158	0.5997
1	−0.6247	−0.625	0.1866	0.5554
2	−0.8784	−0.621	0.1511	0.4922
4	−1.2182	−0.609	0.1125	0.4083
8	−1.6633	−0.588	0.0751	0.3066
16	−2.2102	−0.553	0.0449	0.2038

This is the whole policy. A trader who knows the break-even belief of a setup, the time left to its deadline in information units, and the current belief reads the exit boundary from the second column: exit when $\operatorname{logit}\pi - \operatorname{logit}\theta < g(u)$. The last two columns only illustrate the conversion. The ratio in the third column is almost constant through the swing-trading range and drifts only at the longest deadlines, which is why a one-parameter rule of thumb, $g(u) \approx -0.61\sqrt{u}$, loses little.

The filtered belief against its explicit formula, path by path. *Method.* On the same noise, the belief is computed two ways: by stepping Wonham's equation of Lemma 1 forward with an Euler scheme *in the belief itself*, and by the explicit formula of Proposition 1. Stepping in log-odds would reproduce the formula exactly by construction and would test nothing, so the equation is stepped in the variable in which it was derived. The largest gap over all paths and all times is recorded at three step sizes.

Parameters. $\theta = 0.3$, $p_0 = 0.5$, $S = 1$, 2000 paths; $\delta s = 10^{-2}$, 10^{-3} and 10^{-4}.

Result. The largest gap between the two beliefs is 0.0204, 0.0056 and 0.0023.

Error. The gap shrinks as the step shrinks, by factors of 3.7 and 2.4 per tenfold refinement, close to the $\sqrt{10} \approx 3.2$ of the Euler scheme's strong order one half. The residual is the Euler scheme's error in a multiplicative noise near the edges of $[0, 1]$, not a disagreement between the lemma and the proposition, which agree in the limit.

Patience across the deadline. *Method.* Algorithm 1 at $\theta = 0.3$ for five deadlines, reading the boundary at entry, at mid-life and one step before the deadline. Proposition 4 makes two predictions: patience at entry grows with the deadline, and the boundary depends on the time left only, so that the boundary at the middle of a deadline equals the boundary at entry for half that deadline.

Parameters. $S = 0.25$, 0.5, 1, 2 and 4; $N = 4801$; $\delta s = \min(10^{-4}, S/4000)$.

Result. Patience at entry, $\theta - b(0)$, is 0.0620, 0.0847, 0.1137, 0.1492 and 0.1879. The boundary at mid-life of each deadline equals the boundary at entry of the next shorter one: $b(\frac{1}{2}; 1) = 0.21528 = b(0; \frac{1}{2})$, $b(1; 2) = 0.18628 = b(0; 1)$ and $b(2; 4) = 0.15076 = b(0; 2)$.

One step before the deadline the boundary is 0.29783 for every S, and at the deadline it is 0.3.

Error. The mid-life identities hold to all five printed digits, which is the time-homogeneity of Proposition 4 verified by solves that do not share a time grid. The value of the derived rule against holding to the deadline grows with the runway: 0.0501 against 0.05 at $S = 0.25$, 0.9443 against 0.8 at $S = 4$.

Script. `fig_value.py` — Algorithm 1 at $\theta = 0.3$, $S = 1$, $N = 9601$, $\delta s = 5 \times 10^{-5}$, with the value stored at 201 times.

Shows. Left, the value of holding on over time and belief, from white where holding is worth nothing to navy where it is worth most, with the exit boundary in gold rising from 0.186 at entry to $\theta = 0.3$ at the deadline; the dotted line is the break-even belief the myopic trader uses. Right, three slices of the value against the belief: at entry, at mid-life and near the deadline. Each meets zero smoothly at its boundary point, marked, and the dashed line is the myopic value, which is zero up to θ and has a kink there.

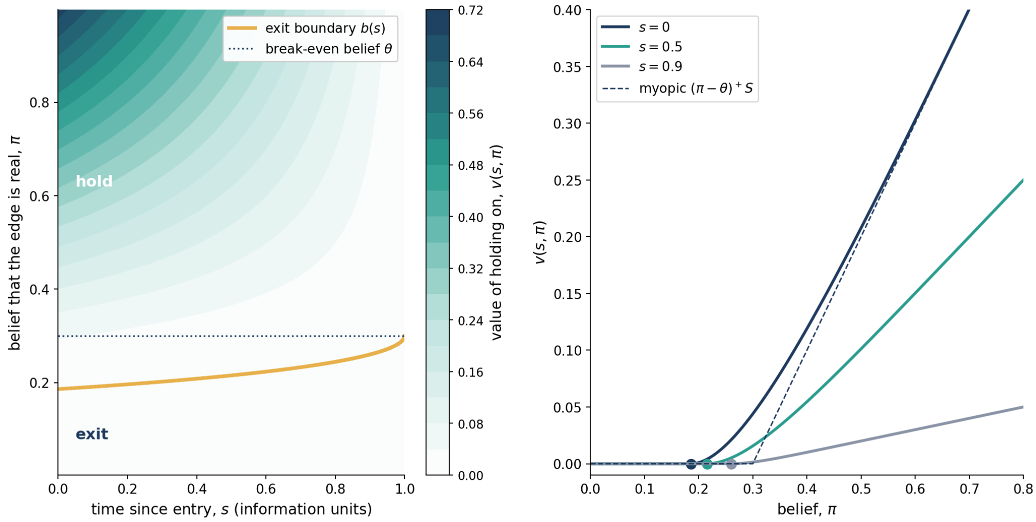

Figure 2: The value of holding on over time and belief, with the exit boundary rising to the break-even belief exactly at the deadline, and three slices showing smooth fit

Script. `fig_fan.py` — Algorithm 1 for five deadlines at $\theta = 0.3$, and for six break-even beliefs at $S = 4$, $N = 9601$.

Shows. Left, the exit boundary against the time since entry for $S = 0.25$, 0.5, 1, 2 and 4, drawn from navy to teal as the deadline lengthens. Every curve ends at $\theta = 0.3$ and a longer runway starts lower: 0.238 with a quarter of a unit, 0.112 with four. Right, the boundary against the time left for $\theta = 0.1$ to 0.6. Each starts at its own break-even belief at the deadline and falls as more time remains. Theorem 2 says these six curves are one curve in log-odds; Figure 1 shows that they are.

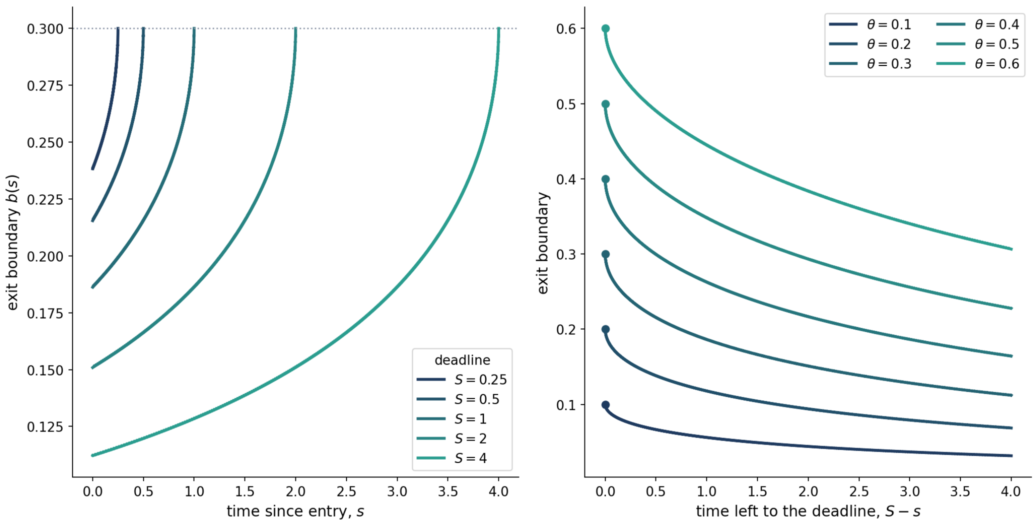

Figure 3: The exit boundary for five deadlines and six break-even beliefs: a longer runway starts lower, every curve ends at its own break-even belief

Script. `fig_price.py` — the boundary of Algorithm 1 at $\theta = 0.3$, $S = 1$, turned into a stop line by Corollary 1 for $p_0 = 0.5$, with fourteen simulated trades of each kind.

Shows. The stop line on the return chart, starting at $x_b(0) = -1.475$ and tightening to $x_b(S) = -0.647$ at the deadline. Left, trades whose edge is real; right, trades with no edge. A red cross marks an exit. Three of the fourteen real edges touch the line and are cut, early and deep; five of the fourteen dead trades are cut, most of them late, when the line has risen to meet them. Most trades of both kinds are held to the deadline, because in one information unit a real edge and no edge are hard to tell apart.

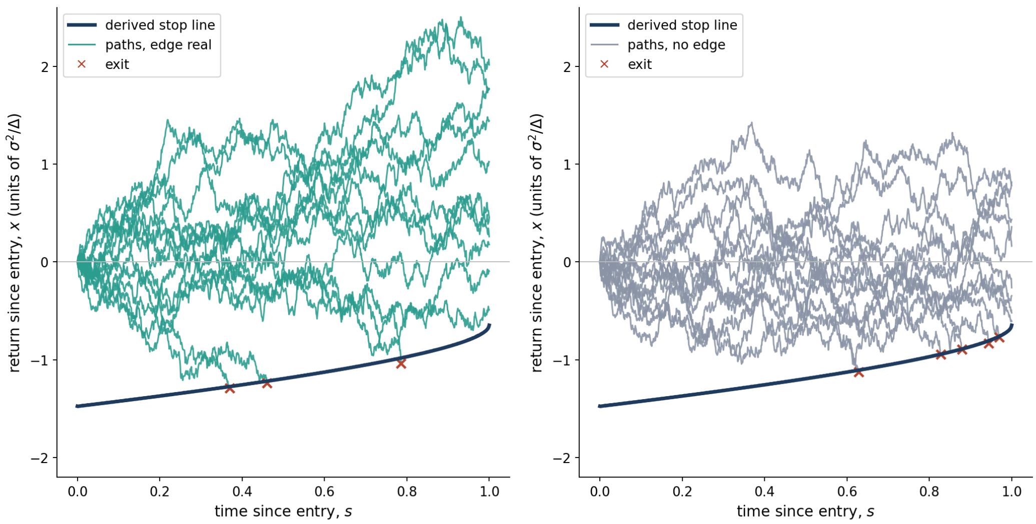

Figure 4: The derived stop line on the return chart, tightening toward the deadline, with simulated trades whose edge is real and trades with no edge

The stop line when the average drift is positive, zero and negative. *Method.* Corollary 1 predicts that the stop line tightens throughout when the average of the two drifts is non-negative, and loosens first when it is negative. The stop line is computed for three break-even beliefs, one on each side of one half and one at it, over a deadline long enough for the loosening to have room.

Parameters. $\theta = 0.3$, 0.5 and 0.7; $S = 4$; $p_0 = 0.8$, above all three entry boundaries; $N = 9601$, $\delta s = 10^{-4}$.

Result. With $\theta = 0.3$ the line rises from -3.453 through -2.714 at mid-life to -1.434 at the deadline. With $\theta = 0.5$ it rises from -2.606 through -2.267 to -1.386. With $\theta = 0.7$ it falls from -1.759 to its lowest point, -1.821, at $s = 1.82$, and then rises to -1.339.

Error. All three shapes are the ones predicted. The turning point for $\theta = 0.7$ is predicted by the square-root approximation of g at $s \approx 1.67$ and measured at 1.82; the difference is the approximation, which slightly overstates $|g'|$ at this range of the time left. At $\theta = 0.5$ the drift part of the line is exactly zero, so the whole rise is the log-odds part, as the corollary says.

Script. `fig_signs.py` — Corollary 1 for $\theta = 0.3$, 0.5 and 0.7 at $S = 4$, $p_0 = 0.8$, from Algorithm 1 at $N = 9601$.

Shows. Left, the three stop lines. When the average drift is positive the line tightens throughout, steeply; when it is zero, less steeply; when it is negative the line first loosens slightly and then tightens toward the deadline. Right, the two parts of each line's move since entry. The log-odds part, dashed, is the same for all three, which is Theorem 2 read on the return chart; the drift part is a straight line whose slope has the sign of the average drift.

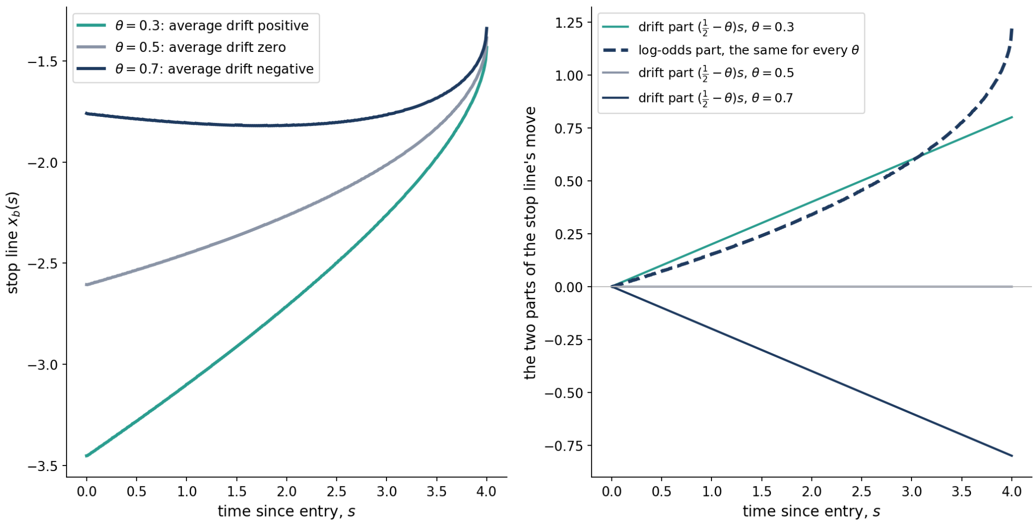

Figure 5: Stop lines that tighten, tighten slowly, or loosen first, depending on the sign of the average drift, and the two parts of their movement

The stop line for five entry beliefs. *Method.* Remark 2 says the entry belief shifts the stop line by $-\operatorname{logit} p_0$ and changes nothing else. The line at entry is computed for five entry beliefs from one solve, and its steps are compared with the steps in log-odds. The value of the derived rule is read at each entry belief and compared with holding to the deadline.

Parameters. $\theta = 0.3$, $S = 1$, $p_0 = 0.3$, 0.4, 0.5, 0.6 and 0.7; the base solve of Algorithm 1.

Result. The stop line at entry is -0.628, -1.070, -1.475, -1.880 and -2.322. Its steps, -0.442, -0.405, -0.405 and -0.442, are the negatives of the steps in $\operatorname{logit} p_0$, which are 0.442, 0.405, 0.405 and 0.442. The value of the derived rule is 0.044, 0.119, 0.208, 0.303 and 0.401; holding to the deadline is worth 0, 0.1, 0.2, 0.3 and 0.4.

Error. The shifts match to the three printed decimals, as they must, because the shift is an identity. The instructive number is the first value. At $p_0 = \theta$ the trade has no expected drift and holding it is worth nothing, yet the derived rule is worth 0.044, all of it the option to learn and leave. As the entry belief rises the option matters less: at $p_0 = 0.7$ the rule adds only 0.001 to holding.

8. When the Rule Exits: the Forward Equation

The HJB equation of the free-boundary derivation runs backward from the deadline and answers one question: *where* should the trader exit? It returns a line. It says nothing about how often a book of trades will meet that line, or when. That is a forward question, about the law of a process that starts at entry and runs toward the deadline, and its equation is the Fokker–Planck equation. The two equations share one object and nothing else: the barrier that the backward equation computes is the absorbing boundary of the forward one.

This section solves the forward equation under each hypothesis. It gives the exit-time density of real edges and of dead trades (Proposition 6), shows that the two densities are tied to each other by the barrier itself (Proposition 7), and reads the result back onto the simulated trades of Figure 4. It ends with what an exit record can and cannot reveal about how many of a trader's setups were real (Remark 3).

Proposition 6 (When the rule exits: a first passage through the HJB boundary, and its forward equation)**.** *Hypotheses.* Theorem 2, Proposition 1 and Corollary 1. *Statement.* In information units the log-likelihood ratio Λ_s starts at 0 and satisfies

$$d\Lambda_s = +\tfrac{1}{2}\, ds + dB_s \quad \text{if the edge is real,} \qquad d\Lambda_s = -\tfrac{1}{2}\, ds + dB_s \quad \text{if it is not.}$$

The derived rule exits at the first passage through the moving barrier

$$\tau^* = \inf\{s \ge 0 :\ \Lambda_s \le z_b(s)\} \wedge S, \qquad z_b(s) = \operatorname{logit}\theta - \operatorname{logit} p_0 + g(S - s).$$

Let $q_H(s, z)$ and $q_L(s, z)$ be the densities of Λ_s on the event that the trade is still open, under each hypothesis. They solve the forward equations

$$\partial_s q_H = -\tfrac{1}{2}\, \partial_z q_H + \tfrac{1}{2}\, \partial_{zz} q_H, \qquad \partial_s q_L = +\tfrac{1}{2}\, \partial_z q_L + \tfrac{1}{2}\, \partial_{zz} q_L, \qquad z > z_b(s),$$

with the absorbing condition $q_H(s, z_b(s)) = q_L(s, z_b(s)) = 0$ and the initial condition

$q_H(0,\cdot) = q_L(0,\cdot) = \delta_0$. The exit-time densities are the fluxes through the barrier,

$$f_H(s) = \tfrac{1}{2}\,\partial_z q_H(s, z_b(s)), \qquad f_L(s) = \tfrac{1}{2}\,\partial_z q_L(s, z_b(s)),$$

and the probabilities of being held to the deadline are $Q_H(S) = \int q_H(S,z)\,dz$ and $Q_L(S) = \int q_L(S,z)\,dz$.

Proof. By Proposition 1 the rule of Theorem 2 exits when $\operatorname{logit} p_0 + \Lambda_s$ *falls to* $\operatorname{logit}\theta + g(S-s)$*, which is* $\Lambda_s \le z_b(s)$*. The law of* Λ *under each drift is the one computed in the derivation after Lemma 1, rescaled to information units. A Brownian motion with constant drift that is killed on a boundary has a sub-probability density that satisfies the Kolmogorov forward equation away from the boundary and vanishes on it [3]; the probability killed per unit time is the flux through the boundary, which the derivation that follows computes.* □

The Fokker-Planck equations, from the diffusion to the flux through the barrier. *Start.* Proposition 6 states the forward equations; this derivation obtains them from the diffusion, fixes the barrier by a change of frame, and shows that exits and survivors account for every trade.

Steps. Under $\mathbb{P}_H$ the generator of Λ is $\mathcal{L}_H\varphi = \frac{1}{2}\varphi' + \frac{1}{2}\varphi''$. For a smooth test function φ that vanishes near the barrier, Itô's formula stopped at τ^* gives $\frac{d}{ds}\mathbb{E}_H[\varphi(\Lambda_{s\wedge\tau^*})] = \mathbb{E}_H[\mathcal{L}_H\varphi(\Lambda_s);\ s < \tau^*]$, and integrating by parts in z moves the generator onto the density as its adjoint:

$$\partial_s q_H = \mathcal{L}_H^* q_H = -\tfrac{1}{2}\,\partial_z q_H + \tfrac{1}{2}\,\partial_{zz} q_H, \qquad \partial_s q_L = \mathcal{L}_L^* q_L = +\tfrac{1}{2}\,\partial_z q_L + \tfrac{1}{2}\,\partial_{zz} q_L. \tag{8.1}$$

The barrier moves, so it is fixed by measuring each open trade by its distance above it in log-odds, $\xi = z - z_b(s)$, and writing $p(s,\xi) = q(s, \xi + z_b(s))$. The chain rule adds the barrier's velocity to the drift:

$$\partial_s p_H = -\Big(\tfrac{1}{2} - z_b'(s)\Big)\partial_\xi p_H + \tfrac{1}{2}\,\partial_{\xi\xi} p_H, \qquad \partial_s p_L = -\Big(-\tfrac{1}{2} - z_b'(s)\Big)\partial_\xi p_L + \tfrac{1}{2}\,\partial_{\xi\xi} p_L, \qquad \xi > 0, \tag{8.2}$$

with $p(s,0) = 0$ and $p(0,\xi) = \delta(\xi + z_b(0))$. The probability current is $J = (\pm\frac{1}{2} - z_b')p - \frac{1}{2}\partial_\xi p$, and at the barrier, where $p = 0$, only the diffusive part survives, so the rate at which trades leave is $f(s) = \frac{1}{2}\,\partial_\xi p(s,0)$. Integrating the forward equation over $\xi > 0$ and using $p(s,0) = 0$ and decay at infinity gives the balance

$$\frac{d}{ds}\int_0^\infty p(s,\xi)\,d\xi = -\tfrac{1}{2}\,\partial_\xi p(s,0) = -f(s), \qquad \int_0^S f(s)\,ds + \int_0^\infty p(S,\xi)\,d\xi = 1. \tag{8.3}$$

The barrier's slope comes from Theorem 2: $z_b'(s) = -g'(S-s)$, positive because g is decreasing, and growing like $(S-s)^{-1/2}$ as the deadline nears. A book in which a share p

of setups carries an edge sees exits at the rate $f(s) = p\, f_H(s) + (1-p)\, f_L(s)$ and holds a share $p\, Q_H(S) + (1-p)\, Q_L(S)$ to the deadline.

End. The drift in the moving frame is the trade's drift *minus the barrier's velocity.* Early in the trade the barrier is almost still and the two hypotheses drift apart at their own rates. Near the deadline the barrier rises steeply, z_b' dominates both, and every trade that is close to the line is swept into it. That is the steep rise of both densities into the deadline in Figure 6. Numerically the moving-frame equations are solved by Crank–Nicolson with the fitted g of the exit-density check.

Proposition 7 (A cut trade had a real edge with probability exactly the exit boundary)**.** *Hypotheses.* Propositions 1 and 6. *Statement.* For every $s < S$ the two exit-time densities satisfy

$$f_H(s) = e^{z_b(s)}\, f_L(s), \qquad z_b(s) = \operatorname{logit} b(s) - \operatorname{logit} p_0.$$

Consequently, in a book in which a share p of setups carries an edge, a trade the rule cuts at time s had a real edge with probability

$$\mathbb{P}(\mu = \mu_H \,|\, \tau^* = s) = \operatorname{expit}\big(\operatorname{logit} p - \operatorname{logit} p_0 + \operatorname{logit} b(s)\big),$$

which equals $b(s)$ exactly when the trader's prior is honest, $p = p_0$.

Proof. The density of $\mathbb{P}_H$ with respect to $\mathbb{P}_L$ on $\mathcal{F}_{\tau^*}$ is $e^{\Lambda_{\tau^*}}$, by optional stopping of the likelihood-ratio martingale at the bounded time τ^*. On the event $\{\tau^* \in ds\}$ with $s < S$ the path is continuous and the rule exits exactly on the barrier, so $\Lambda_{\tau^*} = z_b(s)$. Hence $\mathbb{P}_H(\tau^* \in ds) = \mathbb{E}_L\big[e^{\Lambda_{\tau^*}}; \tau^* \in ds\big] = e^{z_b(s)}\, \mathbb{P}_L(\tau^* \in ds)$. Bayes' rule with prior odds $p/(1-p)$ gives the posterior, and with $p = p_0$ its log-odds is $\operatorname{logit} p_0 + z_b(s) = \operatorname{logit} b(s)$. The proposition says two things about the densities of Figure 6. They have the same shape, because their ratio is set by the barrier, which moves slowly: with $\theta = 0.3$, $S = 1$ and $p_0 = 0.5$ it climbs from $e^{z_b(0)} = 0.230$ at entry to $\theta(1-p_0)/((1-\theta)p_0) = 0.429$ at the deadline. And a late cut is more likely than an early one to have been a real edge, but a cut trade was never a real edge with probability above θ: the rule cuts only trades it believes more likely dead than alive, by exactly the margin the boundary sets. $\square$

The exit-time density from the forward equation against simulation. *Method.* Proposition 6's forward equation is solved by Crank–Nicolson in the moving frame, once under each hypothesis, and compared with simulation. The barrier's slope $z_b'(s)$ needs a smooth g: the numerical boundary of Algorithm 1 carries a grid sawtooth that differentiation would amplify, so g is fitted by $\sum_{k=1}^{4} a_k u^{k/2}$, a series that matches the square-root shape of Table 1. The fit's residual against the computed boundary has root mean square 6.7×10^{-4}, the size of the sawtooth itself, and its leading coefficient is -0.609. The simulation checks the barrier on a grid, and a grid check misses paths that cross and recover between checks, so the barrier is shifted by the continuity correction $0.5826\sqrt{\delta s}$.

Parameters. $\theta = 0.3$, $S = 1$, $p_0 = 0.5$; forward equation on $\xi \in [0, 14]$ with $\delta\xi = 0.002$ and $\delta s = 10^{-4}$, started from the exact Gaussian at $s = 0.02$, when absorption is still below 10^{-20}; 4×10^5 simulated paths per hypothesis at $\delta s = 5 \times 10^{-4}$.

Result. With the edge real the forward equation cuts 13.79 per cent of trades before the deadline and the simulation 13.79 ± 0.06 per cent; with no edge, 44.74 per cent against 44.71 ± 0.08 per cent. Over the four windows $[0.2, 0.4)$, $[0.4, 0.6)$, $[0.6, 0.8)$ and $[0.8, 0.95)$ the two exit probabilities agree to within 0.0005 in every window under both hypotheses. Among trades that are cut, the median exit comes at $s = 0.72$ when the edge is real and $s = 0.68$ when it is not.

Error. The forward solution conserves probability to 10^{-5} with the edge real and 2×10^{-5} without: exits plus survivors sum to one within that. The continuity correction is not optional. Without it the same simulation cuts 13.46 per cent of real edges at this step size and 13.70 per cent at a step five times smaller, converging slowly from below toward the forward equation's 13.79; with it the two agree at the coarser step. Every simulated share quoted in this paper uses the correction.

The two densities against the ratio identity, and the share of real edges among cuts. *Method.* Proposition 7 is tested two ways. First, the two forward equations of Proposition 6 are solved *separately*, one with drift $+\frac{1}{2}$ and one with $-\frac{1}{2}$, and the ratio of their exit-time densities is compared with $e^{z_b(s)}$ at every time step. Nothing in either solve knows the identity. Second, among simulated trades cut in a window of time, the share that had a real edge is compared with the boundary $b(s)$ over the window; with equal numbers of real and dead trades simulated, the population share equals the prior, $p = p_0 = \frac{1}{2}$.

Parameters. $\theta = 0.3$, $S = 1$, $p_0 = 0.5$; the forward solves of the exit-density check; 4×10^5 simulated paths under each hypothesis at $\delta s = 5 \times 10^{-4}$, continuity-corrected.

Result.

s	$f_H(s)$	$f_L(s)$	f_H/f_L	$e^{z_b(s)}$
0.3	0.07495	0.29629	0.252973	0.252971
0.5	0.15713	0.57280	0.274326	0.274325
0.7	0.19477	0.64212	0.303330	0.303329
0.9	0.23909	0.68026	0.351466	0.351464

Over all s between 0.25 and 0.99 the largest relative gap between f_H/f_L and $e^{z_b(s)}$ is 1.7×10^{-5} and the median 3.7×10^{-6}. Among simulated cuts, the share with a real edge is 0.2090 ± 0.0019 for $s \in [0.3, 0.5)$, 0.2234 ± 0.0017 for $[0.5, 0.7)$ and 0.2464 ± 0.0016 for $[0.7, 0.9)$, against the boundary's average over each window, 0.2083, 0.2236 and 0.2450.

Error. The identity holds in the two forward solves to the accuracy of the solves themselves, five orders of magnitude below the size of the densities. Both solves use the same fitted barrier, so this checks the forward solver and the identity, not the fit of g; the simulation, which uses the numerical boundary directly, checks the rest, and all three windows agree with $b(s)$ within one standard error.

Script. `fig_exit_density.py` — the forward equation of Proposition 6 under each hypothesis, with 4×10^5 continuity-corrected simulated paths per panel behind it, at $\theta = 0.3$,

$S = 1$, $p_0 = 0.5$.

Shows. When the derived rule of Figure 4 exits. Left, trades whose edge is real; right, trades with no edge. The smooth curve is the exit-time density from the forward equation and the grey bars are the simulation: they agree across the whole life of the trade. No trade is cut before $s \approx 0.12$, because the stop line starts far below the entry. Dead trades are then cut at a steady rate from about $s = 0.5$, three times as often as real ones. Both densities rise steeply into the deadline, where the stop line of Corollary 1 climbs fastest and catches the trades hovering just above it. The rest of the mass, 86.2 per cent of the real edges and 55.3 per cent of the dead trades, is held to the deadline.

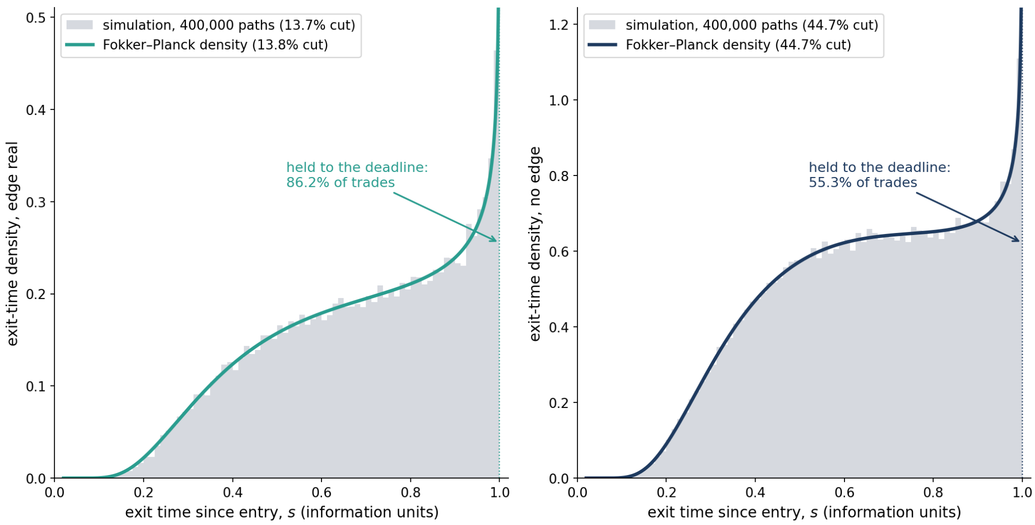

Figure 6: When the derived rule exits: the exit-time density from the forward Fokker–Planck equation through the HJB boundary, for real edges and for dead trades, against simulation

Figure 6 is what Figure 4 becomes when its fourteen paths per panel become four hundred thousand. Every red cross in Figure 4 is a single draw of the exit time τ^*; the density is the law those draws come from. Three of Figure 4's fourteen real edges are cut, 21 per cent, and five of its fourteen dead trades, 36 per cent, against 13.8 and 44.7 per cent in Figure 6. With fourteen paths the binomial standard deviation of those shares is about 0.09 and 0.13, so the small sample is exactly as far from the density as chance allows. The dead trades' crosses in Figure 4 bunch late, four of five after $s = 0.8$, which is the steep rise of the no-edge density into the deadline.

The link runs deeper than sampling. The forward equation of Proposition 6 evolves the *whole cloud* of Figure 4's paths at once: $p(s, \xi)$ is the density of the trades still open at time s, measured by how far above the stop line they are, in log-odds. Each path of Figure 4 is one trajectory of that cloud, and the cloud's leakage through the line is Figure 6. And every cross in Figure 4 sits exactly on the stop line, because a trade is cut the moment it touches it. The return at an exit is therefore not random once the exit time is known: it is $x_b(\tau^*)$, between -1.475 at entry and -0.647 at the deadline. Figure 6's time axis is also a loss axis. An early cut is a deep loss, a late cut a shallow one, and the narrow band of exit returns in Figure 9 is the image of Figure 6 under the stop line.

Remark 3 (Measuring the share of setups with a real edge from the exit record). A trader runs the derived rule, set with prior p_0, on N setups of the same kind. The true share of them that carry an edge, p, is unknown, and it need not equal p_0. Each trade leaves either a cut time s_i or survival to the deadline, and by Proposition 6 the likelihood of the record is

$$\mathcal{L}(p) = \prod_{i \text{ cut}} \Big[p\, f_H(s_i) + (1-p)\, f_L(s_i)\Big] \times \Big[p\, Q_H(S) + (1-p)\, Q_L(S)\Big]^{N_{\text{held}}},$$

maximised at the root $\hat{p}$ of its score. The simplest alternative ignores the timing and uses only the share of trades cut, $\hat{c}$:

$$\hat{p}_{\text{count}} = \frac{c_L - \hat{c}}{c_L - c_H}, \qquad c_H = \int_0^S f_H = 0.138, \qquad c_L = \int_0^S f_L = 0.447.$$

How much does the timing add? The Fisher information per trade is $\int_0^S \frac{(f_H - f_L)^2}{p f_H + (1-p) f_L}\, ds + \frac{(Q_H - Q_L)^2}{p Q_H + (1-p) Q_L}$ with the timing, and $\frac{(c_H - c_L)^2}{c(1-c)}$ with only the count, $c = pc_H + (1-p)c_L$. At $p = 0.5$ they are 0.4657 and 0.4628: the timing adds 0.6 per cent, and 0.4 and 1.1 per cent at $p = 0.3$ and 0.7. On 4000 simulated books of 100 trades with $p = 0.5$, the full estimator has mean 0.502 and standard deviation 0.145, the count estimator mean 0.503 and standard deviation 0.145, matching the Fisher prediction of 0.147.

So the exit record does measure how many of a trader's setups were real, without bias, but almost all of the information is in *how many* trades were cut, not *when*. Proposition 7 explains why: the two densities differ in height far more than in shape, since their ratio only moves from 0.230 to 0.429 over the life of a trade. A useful precision needs a long record. One hundred trades pin the share of real edges to about ± 0.15; ± 0.03 takes about 2400.

9. What a Fixed Time Stop Costs

Theorem 3 (A time stop that ignores the price earns exactly the expected drift over the time held)**.** *Hypotheses.* Proposition 2 and Lemma 1. *Statement.* Let $h \in [0, S]$ be an exit time chosen without reference to the price: a fixed day, or any random time independent of the return. Then, in information units,

$$\mathbb{E} \int_0^h (\pi_s - \theta)\, ds = (p_0 - \theta)\, \mathbb{E}[h].$$

The best such rule is to hold to the deadline if $p_0 > \theta$ and not to enter otherwise. No time stop that ignores the price is ever better than one of those two.

Proof. The belief is a martingale by Lemma 1, so $\mathbb{E}[\pi_s] = p_0$ *for every* s*, and* h *is independent of the path. Exchanging expectation and integral gives* $(p_0 - \theta)\, \mathbb{E}[h]$*, which is linear in* $\mathbb{E}[h]$ *and therefore maximised at an end of* $[0, S]$*. The theorem is short and its* *consequence is not. A time stop's only possible value comes from what the price has done by the time it fires. A trader who closes every position on day three whatever it has done has not managed risk; that trader has shortened every trade by the same amount and earns exactly the expected drift over the shorter holding. The conditional stop of Theorem 4 is the version that can add value, and it is the version traders actually use.* □

Theorem 4 (The best rule that looks once, at day h, in closed form)**.** *Hypotheses.* Propositions 1 and 2, Theorem 3. *Statement.* Among rules that hold to day $h \in (0, S)$, observe the return once, and then either exit or hold to the deadline, the best continues exactly when $\pi_h > \theta$, that is when the log-likelihood ratio exceeds $k^* = \operatorname{logit}\theta - \operatorname{logit} p_0$. A rule that continues when the ratio exceeds a threshold k is worth

$$w(h,k) = (p_0 - \theta)\, h + (S-h)\Big[p_0(1-\theta)\,\Phi\Big(\frac{h/2-k}{\sqrt{h}}\Big) - (1-p_0)\,\theta\,\Phi\Big(\frac{-h/2-k}{\sqrt{h}}\Big)\Big].$$

The trader's rule, continue only if the trade has worked by day h, is $x_h > 0$, which is the threshold $k = -(\frac{1}{2} - \theta)h$.

Proof. Up to day h the rule ignores the price, so Theorem 3 gives the first term. After the look, holding to the deadline from belief π_h *is worth* $(\pi_h - \theta)(S-h)$ *by the same martingale argument, so the best decision is to continue exactly when* $\pi_h > \theta$*. By Proposition 1 that is a threshold on the log-likelihood ratio* Λ_h*, which is normal with mean* $+h/2$ *and variance* h *if the edge is real and mean* $-h/2$ *if it is not. The derivation that follows evaluates* $\mathbb{E}[(\pi_h - \theta)\mathbf{1}\{\Lambda_h > k\}]$ *by conditioning on the truth. The trader's rule follows from* $\Lambda_h = x_h - (\frac{1}{2} - \theta)h$*.* □

The one-look value, from the mixture law of the return. *Start.* Theorem 4's closed form comes from splitting the expected continuation value by the truth, the same move that proves Theorem 2.

Steps. On the event that the rule continues, the value after the look is $(\pi_h - \theta)(S-h)$. Because $\pi_h = \mathbb{P}(\mu = \mu_H \mid \mathcal{F}_h^X)$ and the event is $\mathcal{F}_h^X$-measurable,

$$\mathbb{E}[\pi_h\,\mathbf{1}\{\Lambda_h > k\}] = p_0\,\mathbb{P}_H(\Lambda_h > k), \qquad \mathbb{E}[\mathbf{1}\{\Lambda_h > k\}] = p_0\,\mathbb{P}_H(\Lambda_h > k) + (1-p_0)\,\mathbb{P}_L(\Lambda_h > k). \tag{9.1}$$

Subtracting θ times the second from the first gives

$$\mathbb{E}\big[(\pi_h - \theta)\mathbf{1}\{\Lambda_h > k\}\big] = p_0(1-\theta)\,\mathbb{P}_H(\Lambda_h > k) - (1-p_0)\,\theta\,\mathbb{P}_L(\Lambda_h > k), \tag{9.2}$$

and under the two drifts Λ_h is normal with mean $\pm h/2$ and variance h, so $\mathbb{P}_H(\Lambda_h > k) = \Phi((h/2-k)/\sqrt{h})$ and $\mathbb{P}_L(\Lambda_h > k) = \Phi((-h/2-k)/\sqrt{h})$.

End. The same formula prices every one-look rule, the optimal one and the trader's, and only the threshold differs. The trader's threshold, $-(\frac{1}{2} - \theta)h$, is set by a quantity with no connection to the break-even point: it asks whether the return is positive, where the optimal rule asks whether the belief is above θ.

One look against simulation, and against watching every day. *Method.* Theorem 4's closed form is checked against Algorithm 2 at four look days, for both the optimal threshold and the trader's. Theorem 3 and Theorem 4 bound the derived rule from below, so the value of n evenly spaced looks is also computed, by backward induction on the log-odds grid with Gauss–Hermite quadrature, and must rise toward the PDE value as n grows.

Parameters. $\theta = 0.3$, $S = 1$, $p_0 = 0.5$; look days $h = 0.125$, 0.25, 0.5 and 0.75; 10^6 paths at $\delta s = 10^{-3}$; n from 1 to 128 looks on a grid of 6401 nodes with 96 quadrature points.

Result.

look day h	**optimal, closed form**	**optimal, simulated**	**trader's, closed form**	**trader's, simulated**
0.125	0.20019	0.20002 ± 0.00010	0.14797	0.14855 ± 0.00071
0.25	0.20156	0.20139 ± 0.00021	0.16763	0.16816 ± 0.00064
0.5	0.20437	0.20435 ± 0.00027	0.18950	0.19032 ± 0.00050
0.75	0.20406	0.20404 ± 0.00022	0.19878	0.19929 ± 0.00034

The best single look is at $h = 0.6$ and is worth 0.20470. With n looks the value is 0.20000, 0.20436, 0.20607, 0.20693, 0.20732, 0.20750, 0.20759 and 0.20764 for $n = 1$, 2, 4, 8, 16, 32, 64 and 128, against the PDE value 0.20767.

Error. Every simulated value is within 1.7 standard errors of its closed form; the four trader's rows share paths, so their deviations are correlated and count as fewer than four independent checks. The ladder of looks rises monotonically to the PDE value and stops 3×10^{-5} short of it at 128 looks, as it must, since watching continuously can never be worth less. An earlier run on a coarser grid put the 128-look value 4×10^{-5} *above* the PDE value, which is impossible; refining the look grid removed it, and it is recorded here because a bound that is violated is a bug report, not a result.

Script. `fig_onelook.py` — the n-look induction for n from 1 to 128, Theorem 4's closed form over the look day, and the PDE value from Algorithm 1, at $\theta = 0.3$, $S = 1$, $p_0 = 0.5$.

Shows. Left, the value of looking n times, on a logarithmic axis, rising from holding to the deadline (0.2) toward the derived rule (0.20767): two looks capture more than half of the gain, and the rest comes slowly. Right, the value of a single look against the day it is taken. The optimal threshold, navy, is never worse than holding and is best near $h = 0.6$. The trader's rule, red, is worse than holding at every look day, and far worse when the look is early.

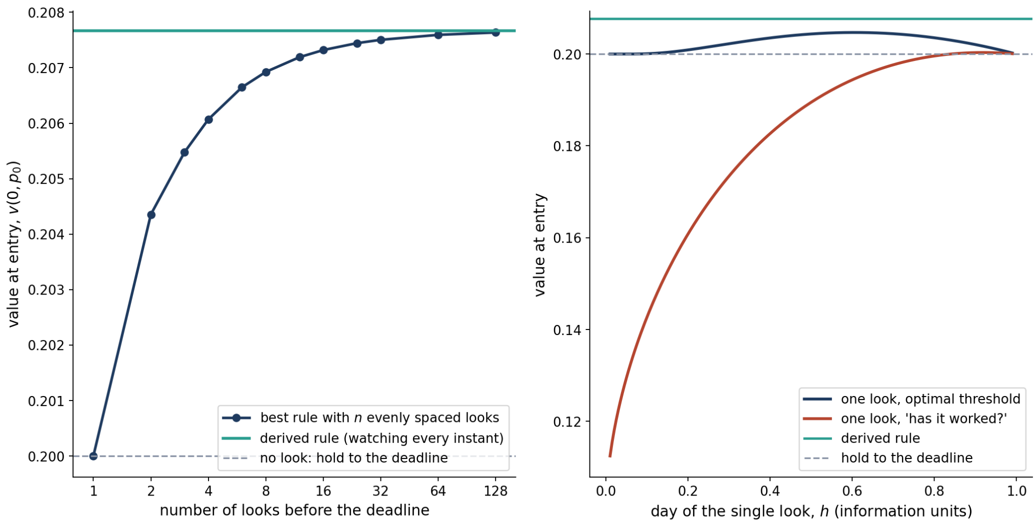

Figure 7: The value of watching: holding without looking, one look at the best day, many looks, and the trader's single look, against the derived rule

Value lost by fixed stops over the two dimensionless numbers. *Method.* Over a grid of break-even beliefs and deadlines, the derived rule's value from Algorithm 1 is compared with holding to the deadline (Theorem 3) and with the trader's stop at a quarter, a half and three quarters of the deadline (Theorem 4's closed form). Each rule's shortfall is reported as a percentage of the derived rule's value.

Parameters. $p_0 = 0.5$; θ from 0.05 to 0.45 in steps of 0.05; $S = 0.25$, 0.5, 1, 2 and 4; forty-five cases; $N = 2401$.

Result. The trader's stop at $S/2$ gives up between 3.3 per cent (at $\theta = 0.1$, $S = 4$) and 20.6 per cent (at $\theta = 0.45$, $S = 4$) of the attainable value. At $S/4$ it gives up 11.5 to 30.0 per cent, and at $3S/4$ it gives up 0.8 to 36.5 per cent. Holding to the deadline gives up nothing measurable in the most favourable case and up to 61.4 per cent when the break-even belief is close to the entry belief and the runway is long. The trader's stop at $S/2$ is *worse than holding to the deadline* in thirty of the forty-five cases, every one of them with $\theta \le 0.4$.

Error. The closed forms are exact and the PDE values carry the error of the convergence study, below 10^{-6} in the value, so every percentage is exact to the digit printed. The claim verified is Theorem 1's optimality in the only form that matters to a trader: no rule on the grid beats the derived one, and the rule traders use is beaten by it everywhere and by doing nothing in two cases out of three.

Script. `fig_cost.py` — Algorithm 1 on a 25×19 grid of break-even beliefs and deadlines, against Theorem 3 and Theorem 4, at $p_0 = 0.5$.

Shows. The value given up against the derived rule, as a percentage, over the break-even belief and the deadline. Left, the trader's stop at $S/2$: it costs at least three per cent everywhere, most when the deadline is short and the edge strong, and it climbs again as θ approaches the entry belief. Right, holding to the deadline: it costs almost nothing when the edge is strong, and a great deal only when the setup is close to break-even at entry. The two maps are nearly complementary, which is the point of Figure 9: the trader's stop cuts hardest exactly where holding was nearly right.

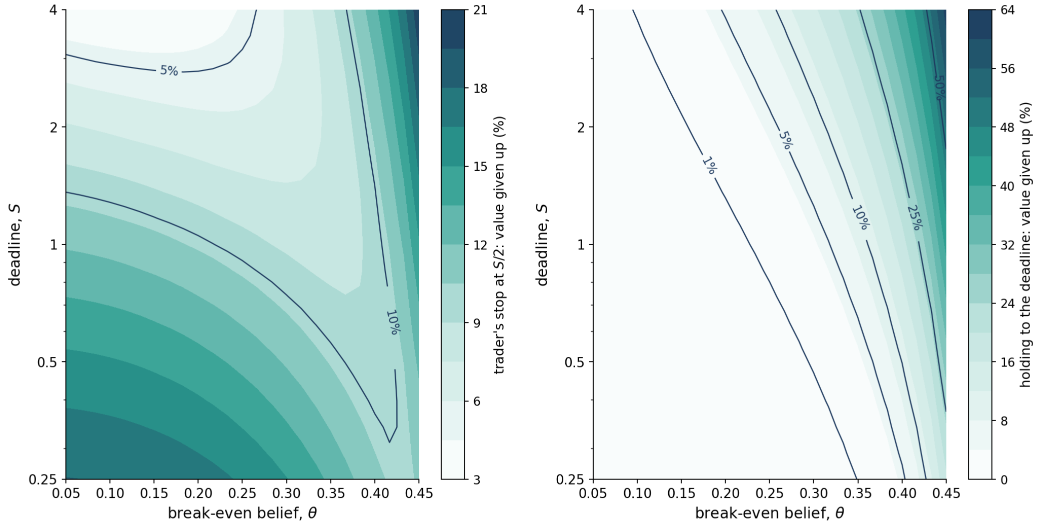

Figure 8: The cost of fixed rules over the break-even belief and the deadline: the trader's stop at half the deadline, and holding to the deadline

Source. `compute_numbers.py`, section 9 — Algorithm 2 on 10^6 paths at $\delta s = 10^{-3}$, $\theta = 0.3$, $S = 1$, $p_0 = 0.5$, every rule on the same paths and estimated paired against holding; the derived rule's boundary carries the continuity correction of Algorithm 2.

Columns. the rule (the derived rule of Theorem 1, holding to the deadline of Theorem 3, and the trader's stop at three look days); its value in units of σ^2/Δ with its standard error; the difference from the derived rule; the mean holding time; the share of trades *with a real edge* cut before the deadline; and the share of trades *with no edge* held to it.

rule	**value**	**vs derived**	**mean hold**	**edges cut**	**dead held**
derived	0.2076 ±0.0003	—	0.904	13.8%	55.3%
hold to S	0.2000	−3.7%	1.000	0.0%	100.0%
stop at $S/4$	0.1682 ±0.0006	−19.0%	0.654	36.3%	44.0%
stop at $S/2$	0.1903 ±0.0005	−8.3%	0.776	31.0%	41.6%
stop at $3S/4$	0.1993 ±0.0003	−4.0%	0.891	27.3%	39.8%

The derived rule cuts one real edge in seven and still carries more than half of the dead trades to the deadline, because in one information unit the two are hard to tell apart and cutting is expensive when it is wrong. The trader's stops cut between a quarter and more than a third of the real edges, which is what they lose, and they do not buy much for it: they still carry about two in five dead trades to the end. An earlier stop cuts more of both kinds, and it is the real edges that carry the value.

Script. `fig_holding.py` — Algorithm 2 on 300 000 paths at $\theta = 0.3$, $S = 1$, $p_0 = 0.5$, the derived rule and the trader's stop at $S/2$ on the same paths.

Shows. Left, the share of trades cut by each moment, separately for trades whose edge is real (solid) and trades with no edge (dashed). The derived rule cuts gradually, and cuts dead trades more than three times as often as real ones, 44 per cent against 14. The trader's stop cuts everything it will ever cut at once, at $S/2$, and cuts 31 per cent of the real edges to 58 per cent of the dead trades. Right, the return at an early exit. The derived rule's exits sit in a narrow band along its stop line, between about -1.50 and -0.64. The trader's exits are spread from deep losses to zero, because a stop that looks only at whether the trade is up ignores how far down it is.

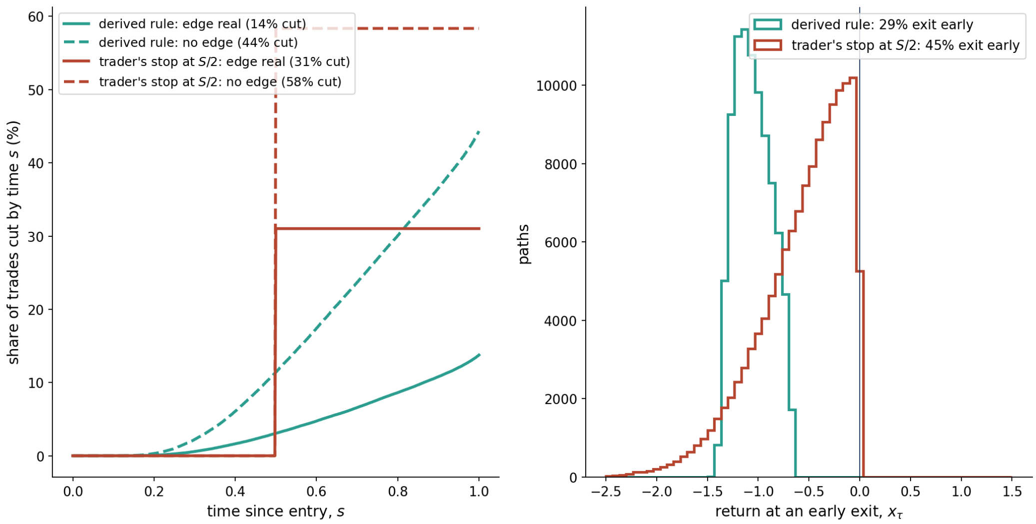

Figure 9: The share of trades cut over time and the return at an early exit, for the derived rule and the trader's stop at half the deadline

The practical content fits in five sentences.

A time stop that ignores the price does nothing but shorten the trade: Theorem 3 prices it at exactly the expected drift over the shorter holding, so it can never beat holding to the deadline. The time stop traders actually use, *out unless it has worked by day h*, asks the wrong question. Whether the trade is up measures the return against zero, while the decision turns on whether the belief is above the break-even point, and the two thresholds agree only by coincidence (Theorem 4). The derived rule is one table: exit when the log-odds that the edge is real fall more than $|g(u)|$ below the log-odds of break-even, with g read from Table 1 at the time left u, and $g(u) \approx -0.61\sqrt{u}$ over the whole swing range. On the chart it is a stop line that starts wide and tightens toward the deadline, and it sits deeper the more conviction the trader had at entry. Without a deadline there is no stop, so a trader who wants one must decide, before entry, when the setup expires.

To use the rule a trader needs three numbers about a setup. The first is the break-even belief $\theta = -\mu_L/\Delta$: how much of the gap between a working and a failed setup is the failed setup's loss. The second is the edge-to-noise ratio Δ/σ, which fixes the size of an information unit, $(\sigma/\Delta)^2$ of calendar time, and of a return unit, σ^2/Δ. The third is an honest entry belief. None of them is assumed anywhere in this paper; every result is stated in the units they define, so a trader who has them can read the answer and a trader who does not can read the shape.

10. Limits

Remark 4 (Two drifts, constant volatility, no costs: what each assumption buys). Each assumption of Assumption 1 buys a specific simplification, and each can be relaxed at a specific price. *Two drifts.* A setup either works or it does not, and the belief is one number. With a continuum of possible drifts the belief becomes a distribution; with a Gaussian prior it is summarised by its mean and a deterministic variance, and the problem is again two-dimensional, but the universality of Theorem 2 is specific to two drifts, because it rests on the value splitting into exactly two holding times. Ekström and Vaicenavicius [6] treat general priors for the selling problem. *Constant, known volatility.* The information unit $(\sigma/\Delta)^2$ is fixed. A volatility that moves makes information arrive faster when the price is calm, since the same drift is easier to see through less noise. *One unit, no costs.* A cost paid at exit makes cutting more expensive and lowers the boundary by an amount that can be computed with the same scheme and a changed obstacle; a cost of carrying the position raises the break-even belief, and Theorem 2 then covers it unchanged. *Expected return.* The criterion suits a trader running many small independent setups. A trader with one large position cares about its variance, and the problem becomes one of risk-sensitive stopping.

None of these limits touches Theorems 3 and 4. A time stop that ignores the price is worth $(p_0 - \theta)$ per unit of time held under any prior and any volatility, because the belief is a martingale whatever it is a belief about.

Remark 5 (A continuum of drifts, a trailing stop, and a scheduled catalyst). Three extensions change the question rather than the numbers. A continuum of drifts turns *does it work* into *how well does it work*, and the exit rule then depends on the whole posterior rather than on one probability. A trailing stop is a boundary in the running maximum of the return rather than in the return itself. The derived stop line of Corollary 1 does not trail: it depends on time and on the entry belief, not on the best level reached. Whether a trailing component adds value under drift uncertainty is a separate question, closer to the problem of selling at the maximum [10]. A scheduled catalyst, an earnings date inside the window, adds a jump of known timing and unknown size. The deadline of this paper then becomes a decision to hold through the event or not, and the boundary would acquire a second terminal condition at the event date.

11. Conclusion

A swing trade with an unknown edge and a known deadline has an exact exit rule. The trader holds while the belief that the edge is real stays above a boundary that rises to the break-even belief exactly at the deadline, and the boundary is one universal curve in log-odds, the same for every break-even point, close to -0.61 times the square root of the time left. On the return chart the rule is a stop line that starts wide, tightens toward the deadline unless the average of the two drifts is negative, and sits deeper the more conviction the trader had at entry. Without a deadline it disappears.

Against it the rules of thumb are expensive. A time stop that ignores the price is worth exactly the expected drift over the time held and never beats holding to the deadline. The conditional time stop traders use gives up three to twenty per cent of the attainable value in the cases computed here, and is worse than simply holding in two cases out of three, because asking whether a trade is up is not the same as asking whether it is still worth holding. It cuts real edges in quantity to save dead trades it mostly keeps anyway.

Every number in this paper is computed in units that assume nothing about a market, from a finite-difference solution checked against closed forms, two bounds and a million-path simulation. A reader with three numbers about a setup, its break-even belief, its edge-to-noise ratio and an honest entry belief, has everything the rule needs.

References

[1] Wonham, W. M. (1964). Some Applications of Stochastic Differential Equations to Optimal Nonlinear Filtering. *Journal of the Society for Industrial and Applied Mathematics Series A Control*. https://doi.org/10.1137/0302028

[2] Liptser, R. S., & Shiryaev, A. N. (2001). *Statistics of Random Processes*. Springer. https://doi.org/10.1007/978-3-662-13043-8

[3] Peskir, G., & Shiryaev, A. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhauser. https://doi.org/10.1007/978-3-7643-7390-0

[4] Gapeev, P. V., & Peskir, G. (2004). The Wiener Sequential Testing Problem with Finite Horizon. *Stochastics and Stochastic Reports*. https://doi.org/10.1080/10451120410001663753

[5] Ekstrom, E., & Lu, B. (2011). Optimal Selling of an Asset under Incomplete Information. *International Journal of Stochastic Analysis*. https://doi.org/10.1155/2011/543590

[6] Ekstrom, E., & Vaicenavicius, J. (2016). Optimal Liquidation of an Asset under Drift Uncertainty. *SIAM Journal on Financial Mathematics*. https://doi.org/10.1137/15M1033265

[7] Shiryaev, A., Xu, Z., & Zhou, X. Y. (2008). Thou Shalt Buy and Hold. *Quantitative Finance*. https://doi.org/10.1080/14697680802563732

[8] Dai, M., Zhang, Q., & Zhu, Q. J. (2010). Trend Following Trading under a Regime Switching Model. *SIAM Journal on Financial Mathematics*. https://doi.org/10.1137/090770552

[9] Decamps, J.-P., Mariotti, T., & Villeneuve, S. (2005). Investment Timing under Incomplete Information. *Mathematics of Operations Research*. https://doi.org/10.1287/moor.1040.0132

[10] du Toit, J., & Peskir, G. (2009). Selling a Stock at the Ultimate Maximum. *The Annals of Applied Probability*. https://doi.org/10.1214/08-AAP566

[11] Kaminski, K. M., & Lo, A. W. (2014). When Do Stop-Loss Rules Stop Losses?. *Journal of Financial Markets*. https://doi.org/10.1016/j.finmar.2013.07.001

[12] Jaillet, P., Lamberton, D., & Lapeyre, B. (1990). Variational Inequalities and the Pricing of American Options. *Acta Applicandae Mathematicae*. https://doi.org/10.1007/BF00047211

[13] Barles, G., & Souganidis, P. E. (1991). Convergence of Approximation Schemes for Fully Nonlinear Second Order Equations. *Asymptotic Analysis*. https://doi.org/10.3233/asy-1991-4305