Switching Cost Models as Hypothesis Tests*

Samuel N. Cohen
University of Oxford

Timo Henckel
Australian National University

Gordon D. Menzies
University of Technology & CAMA†

Johannes Muhle-Karbe
Carnegie Mellon University

Daniel John Zizzo
Newcastle University

July 14, 2018

Abstract

Models based on belief-based switching costs (such as inattention costs) are related to hypothesis tests and vice versa. Specifically, an inference problem with a penalty for mistakes and for switching the inferred value gives a band of inaction which is equivalent to a confidence interval, and therefore to a two-sided hypothesis test.

Keywords: inference; switching cost; inferential expectations.

JEL classification codes: D01, D81, D84

1 Introduction

This paper provides a new micro-foundation for a hypothesis test. Agents receive sequential information and conduct inference which penalizes both the variance of the estimator and adjustments to the estimator whenever it changes. The fully optimal estimator has a band of inaction, whereby it is updated only when the classical Bayesian estimate leaves it. We show that, to a first order approximation for small adjustment costs, this band of inaction has width proportional to the Bayesian estimator's standard deviation, making it equivalent to a confidence interval and therefore to a two-sided hypothesis test.

Our result locates belief formation models based on hypothesis tests, such as [10], within a wider literature on switching costs due to sticky belief adjustment and vice versa. The switching costs may arise, say, from cognitive effort

*We acknowledge helpful comments of colleagues at the University of Technology Sydney.

†Gordon D. Menzies (corresponding author), School of Economics and Finance, University of Technology Sydney, Broadway, Sydney NSW 2007, Australia. E-mail address: gordon.menzies@uts.edu.au. Phone: +61-2-95147728.

in attention and observation, information gathering, or the consultation of experts [3, 9, 5]. State-dependant belief adjustment implies a dependence on the economic state, which in turn depends on new information arriving. Models of inattention and portfolio choice are example applications of this approach [1, 7].

The theory of optimal investment under transaction costs has been well studied within the mathematical finance literature. A key approach, which our approach is based on, is the use of asymptotic approximation methods, to allow closed-form solutions which are valid when transaction costs are small. For both proportional costs [11] and fixed costs [2, 6], one is able to derive an approximate 'no trade region', within which agents accept deviations from the no-transaction-cost optimal strategy, in order to avoid excessive costs. In this paper, we translate this analysis to an estimation setting, giving the desired bands of inaction. Our approach follows [2] closely, however their analysis is not directly applicable, as our setting is not based on trading in a financial market.

2 The model

We base our setting on a Kalman–Bucy filter [8], as this gives a wide range of applications.

We write X_t for a multivariate (hidden) process, which we seek to estimate using multivariate observations Y. We suppose X and Y satisfy

$$\begin{aligned} dX_t &= F_t X_t dt + dW_t; \qquad X_0 \sim N(\hat{X}_0, P_0) \\ dY_t &= H_t X_t dt + dB_t \end{aligned}$$

where W and B are independent continuous martingales, with quadratic variations

$$d\langle W \rangle_t = Q_t dt, \qquad d\langle B \rangle_t = R_t dt.$$

Here F, H, Q and R are matrix-valued deterministic processes of appropriate dimensions, we assume R is invertible and H is nonzero, and $\hat{X}_0, P_0$ are known. We define a filtration by $\mathcal{F}_t = \sigma(Y_s; s \leq t)$, which represents the information available from observing Y up to time t.

The key result of [8] is that, conditional on our observations, X_t has a multivariate normal distribution, that is,

$$X_t | \mathcal{F}_t \sim N(\hat{X}_t, P_t).$$

The values of $(\hat{X}_t, P_t)$ have dynamics

$$\begin{aligned} d\hat{X}_t &= F_t \hat{X}_t dt + K_t d\hat{V}_t, \\ dP_t &= F_t P_t + P_t F_t^\top + Q_t - K_t R_t K_t^\top, \end{aligned}$$

with initial values (X_0, P_0), where $K_t = P_t H_t^\top R_t^{-1}$ denotes the *Kalman gain* process, and $d\hat{V}_t = dY_t - H_t \hat{X}_t dt$ defines the *innovations process* $\hat{V}$, which is a (multivariate) Brownian motion under $\{\mathcal{F}_t\}_{t \geq 0}$.

Example 1. A simple example is when our processes are all scalar, $F, Q \equiv 0$ and $H, R \equiv 1$. Then $X = X_0$ is a constant (unknown) quantity and $K_t = P_t$, so

$$\frac{dP_t}{dt} = -K_t R_t K_t^\top = -P_t^2, \qquad \Rightarrow \qquad P_t = \frac{1}{1/P_0 + t},$$
$$d\hat{X}_t = K_t d\hat{V}_t = \frac{1}{1/P_0 + t} d\hat{V}_t.$$

We suppose that, over a fixed time period $[0, T]$, our agent has wealth Z_t, and pays costs $\rho(\hat{X}_t - \theta_t)$ due to tracking error and a cost λ whenever θ_t changes. We assume ρ is smooth and minimized at $\rho(0) = 0$. For a utility function U, our agent wishes to optimize her expected wealth

$$J(\omega, t, z, \theta; \lambda) = E\Big[U\Big(z - \int_t^T \rho\big(\hat{X}_t - \theta_t\big) - \lambda \sum_{t \leq s \leq T} I_{\{\Delta\theta_t \neq 0\}}\Big)\Big|\mathcal{F}_t\Big].$$

over piecewise constant adapted processes θ. As $\hat{X}$ is a Markov process, there exists a value function

$$v(t, \hat{X}_t, z, \theta_t; \lambda) = \sup_{\theta' : \theta_t = \theta'_t} J(\omega, t, z, \theta'; \lambda).$$

The approach of [2] can now be employed[1]. We expand v in terms of powers of λ and determine the corresponding coefficients. By ignoring higher order terms, we obtain a first approximation to our value function, and hence to the optimal choice of θ.

Given the presence of a fixed cost, the optimal policy will be to leave θ unchanged until $\hat{X}_t - \theta_t$ is large. Write $\mathfrak{K}$ for the region where θ remains fixed. A standard dynamic programming argument yields a partial differential equation for v:

In the regime where θ_t does not change, $\hat{X}_t$ follows a diffusion process. The martingale principle of optimality[2], combined with Itô's lemma, shows that v satisfies the inequality

$$0 \geq \partial_t v - (\partial_z v) u\big(\hat{X}_t - \theta_t\big) + (\partial_{\hat{x}} v)^\top F \hat{X}_t + \frac{1}{2}\mathrm{Tr}(K_t K_t^\top \partial_{\hat{x}\hat{x}} v), \tag{1}$$

with equality on $\mathfrak{K}$ (when it is not optimal to change θ). Considering the possibility of changing θ, we observe

$$v(t, z, \hat{x}, \theta; \lambda) \geq \sup_{\theta'} v(t, z - \lambda, \hat{x}, \theta'; \lambda) \tag{2}$$

[1]In [2] this analysis is completed, as the required regularity assumptions on v are checked, to ensure that we have obtained an approximately optimal strategy. (See also [6] for a different approach.) The corresponding calculations are lengthy and lend little to our qualitative understanding.

[2]This is a form of the dynamic programming principle, and states that the value function should be a martingale under the optimal strategy, and a supermartingale otherwise.

with equality on $\mathfrak{K}^c$. Combining these two inequalities, we obtain the dynamic programming equation

$$0 = \min \Big\{ -\partial_t v + (\partial_z v) u\big(\hat{x} - \theta\big) - (\partial_{\hat{x}} v)^\top F \hat{x} - \frac{1}{2}\mathrm{Tr}(K_t K_t^\top \partial_{\hat{x}\hat{x}} v), \\ v(t, z, \hat{x}, \theta; \lambda) - \sup_{\theta'} v(t, z - \lambda, \hat{x}, \theta'; \lambda) \Big\},$$

with terminal value $v(T, z, \hat{x}, \theta; \lambda) = U(z)$. The difficulty in solving this equation directly is that we have a free boundary for $\mathfrak{K}$. We focus on obtaining an approximation for small values of λ.

2.1 Asymptotic analysis

When $\lambda = 0$ we know $v(t, z, \hat{x}; 0) \equiv U(z)$. Following [2], we expect that the optimal strategy[3] will involve switching whenever $\|\hat{X} - \theta\| = O(\lambda^{1/4})$, resulting in a cost of $O(\lambda^{1/2})$. This gives the ansatz

$$v(t, z, \hat{x}, \theta; \lambda) = U(z) + \lambda^{1/2}\phi(t, z) + \lambda\psi(t, z, \hat{x}, \xi) + O(\lambda^{3/2}) \tag{3}$$

where $\xi := (\hat{x} - \theta)\lambda^{-1/4}$.

Recalling our assumptions on ρ,

$$\rho(\xi\lambda^{1/4}) = \lambda^{1/2}\xi^\top \Gamma \xi + o(\lambda^{1/2}),$$

where $\Gamma = -\partial_{xx}\rho/2$ is a positive-definite matrix.

We plug-in the ansatz (3) into (1), to obtain (on $\mathfrak{K}$)

$$0 = \lambda^{1/2}\Big(\phi_t + \xi^\top \Gamma \xi U' + \frac{1}{2}\mathrm{Tr}\big(K_t K_t^\top \psi_{\xi\xi}\big)\Big) + o(\lambda^{1/2}).$$

On $\mathfrak{K}^c$, from (2) we have

$$\begin{aligned} 0 &= v(t, z, \hat{x}, \theta; \lambda) - \sup_{\theta'} v(t, z - \lambda, \hat{x}, \theta'; \lambda) \\ &= -\lambda U' - \lambda^{3/2}\phi_z + \lambda\Big(\psi(t, z, \hat{x}, \xi) - \sup_{\xi'} \psi(t, z - \lambda, \hat{x}, \xi')\Big). \end{aligned} \tag{4}$$

We can assume[4] $\psi(t, z - \lambda, \hat{x}, \xi) = \psi(t, z, \hat{x}, \xi) + o(\lambda)$, where the error depends only on z and, as λ multiplies ψ in (3),

$$\sup_{\xi'} \psi(t, z, \hat{x}, \xi') = \psi(t, z, \hat{x}, 0) = 0,$$

which simplifies (4):

$$0 = -\lambda\Big(U' + \psi(t, z, \hat{x}, \xi)\Big) + o(\lambda).$$

[3]This strategy comes from analyzing, over long horizons, how often the boundary of an interval will be hit by a random walk, averaging out the cost paid, then optimizing over the width of the interval chosen. The arguments of [2] hold in our setting, mutatis mutandis.

[4]This assumption can be verified in many cases, see [2].

Considering the leading λ-order term in each region, we obtain

$$\begin{cases} 0 = \phi_t + (\xi^\top \Gamma \xi) U' + \frac{1}{2} \mathrm{Tr}\big(K_t K_t^\top \psi_{\xi\xi}\big) & \text{on } \mathfrak{K}, \\ 0 = U'(z) + \psi(t, z, \hat{x}, \xi) & \text{on } \mathfrak{K}^c. \end{cases}$$

2.2 Exponential Utility

To obtain a closed-form solution, we shall assume that $U(z) = (1 - e^{-kz})/k$ for some $k > 0$. Then $U'(z) = e^{-kz}$ and with $\tilde{\phi} = e^{kz}\phi$, $\tilde{\psi} = e^{kz}\psi$,

$$\begin{cases} 0 = \tilde{\phi}_t + (\xi^\top \Gamma \xi) k + \frac{1}{2} \mathrm{Tr}\big(K_t K_t^\top \tilde{\psi}_{\xi\xi}\big) & \text{on } \mathfrak{K}, \\ 0 = -1 + \tilde{\psi}(t, z, \hat{x}, \xi) & \text{on } \mathfrak{K}^c. \end{cases}$$

In addition to this, a smooth pasting property should hold on the boundary.

Following [4, 2], we propose[5] a solution of the form

$$\tilde{\psi}^*(t, z, \hat{x}, \xi) = -1 + (\xi^\top M \xi - 1)^2$$

with $\mathfrak{K} = \{\xi : \xi^\top M \xi < 1\}$, where M is a (symmetric, positive-definite) matrix to be determined. Calculating,

$$\tilde{\psi}^*_{\xi\xi} = 4(\xi^\top M \xi - 1) M - 8 M \xi \xi^\top M.$$

Therefore,

$$\begin{aligned} 0 &= \tilde{\phi}^*_t + \xi^\top \Gamma \xi + \frac{1}{2} \mathrm{Tr}\big(K_t K_t^\top 4(\xi^\top M \xi - 1) M - 8 M \xi \xi^\top M\big) \\ &= \tilde{\phi}^*_t - 2\mathrm{Tr}\big(K_t K_t^\top M\big) + \xi^\top \Big(\Gamma + 2M \mathrm{Tr}\big(K_t K_t^\top M\big) - 4 M K_t K_t^\top M\Big)\xi. \end{aligned}$$

This should hold for all ξ, so by comparing coefficients

$$0 = \Gamma + 2M \mathrm{Tr}\big(K_t K_t^\top M\big) - 4 M K_t K_t^\top M \tag{5}$$

which is an algebraic equation to be solved for M.

2.3 The scalar setting

In one dimension, (5) simplifies to

$$0 = \Gamma - 2M^2 K_t^2 \qquad \Leftrightarrow \qquad M = \frac{\sqrt{\Gamma/2}}{K_t}.$$

The no-switching region is

$$\mathfrak{K} = \Big\{\xi : \xi \le \frac{\sqrt{K_t}}{(\Gamma/2)^{1/4}} \propto \sqrt{K_t}\Big\}.$$

[5] In one dimension, this is the smallest family of polynomials which could satisfy our assumptions.

In the special case where the hidden variable X_t is constant,

$$P_t = K_t = (1/P_0 + t)^{-1},$$

and as P_t is the variance of $X_t|\mathcal{F}_t$, we observe the desired test-statistic behavior; we switch whenever

$$\frac{|\hat{X}_t - \theta_t|}{P_t^{1/2}} > c$$

for $c = (2\lambda/\Gamma)^{1/4}$. By choosing the test size α of our two-sided test such that c is the critical value of the test statistic, our optimal policy is to switch whenever a hypothesis test is failed.

Remark 1. In the scalar setting, when X is not a constant, it is interesting that the optimal switching does not correspond to a hypothesis test *with fixed test size*. In general, $\mathfrak{K}$ has width $\propto K_t^{1/2} = (P_t H_t/R_t)^{1/2}$. As H_t/R_t describes the quality of observations, and hence the volatility of $\hat{X}$, we see that in periods of low-quality data (so $\hat{X}$ is not volatile) our agent will switch more frequently, as deviations between θ_t and X_t are likely to persist for longer.

References

[1] Andrew B. Abel, Janice C. Eberly, and Stavros Panageas. Optimal inattention to the stock market with information costs and transaction costs. *Econometrica*, 81(4):1455–1481, 2013.

[2] Albert Altarovici, Johannes Muhle-Karbe, and H. Mete Soner. Asymptotics for fixed transaction costs. *Finance and Stochastics*, 19(2):363–414, April 2015.

[3] Fernando Alvarez, Francesco Lippi, and Juan Passadore. Are state and time dependent models really different? *NBER Macroeconomics Annual*, 31, 2016.

[4] C. Atkinson and P. Wilmott. Portfolio management with transaction costs: an asymptotic analysis of the morton and pliska model. *Mathematical Finance*, 5(4):357–367, 1995.

[5] Christopher Carroll. Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, 118:269–298, 2003.

[6] Mark-Roman Feodoria. *Optimal investment and utility indifference pricing in the presence of small fixed transaction costs*. PhD thesis, Christian-Albrechts-Universität zu Kiel, 2016. available at `https://macau.uni-kiel.de/servlets/MCRFileNodeServlet/dissertation_derivate_00006786/Feodoria_Dissertation_Druckversion.pdf`.

[7] Lixin Huang and Hong Liu. Rational inattention and portfolio selection. *Journal of Finance*, 62(4):1999–2040, 2007.

[8] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 1961.

[9] Jacopo Magnani, Aspen Gorry, and Ryan Oprea. Time and state dependence in an ss decision experiment. *American Economic Journal: Macroeconomics*, 8(1):285–310, 2016.

[10] G. Menzies and D. Zizzo. Inferential expectations. *B. E. Journal of Macroeconomics*, 9:1–25, 2009.

[11] Johannes Muhle-Karbe, H. Mete Soner, and Max Reppen. A primer on portfolio choice with small transaction costs. *Annual Review of Financial Economics*, 9:301–331, 2017.

Appendix: Illustration in Discrete Time

In order to give a concrete example, and to justify the asymptotica approximation used, we consider the similar problem of estimating a parameter $p \in [0,1]$, sampling from Bernoulli trials Y_i in order to conduct a hypothesis test of

$$H_0 : p = p_0 \qquad \text{vs.} \qquad H_1 : p \neq p_0.$$

For notational simplicity, we write $\mathcal{F}_t$ for the information available from the first t observations, that is $Y_1, ..., Y_t$.

If the sample size t is large we apply the Central Limit Theorem to the Maximum Likelihood Estimator (MLE) $\hat{p}_t$.

$$\begin{aligned} Y_i &\sim i.i.d.\,(p, p(1-p))\,, \\ \frac{\sum Y_i}{t} &= \hat{p}_t \sim N\left(p, \frac{p(1-p)}{t}\right). \end{aligned}$$

A two-sided hypothesis test of size α is a rule whereby we maintain H_0 if $\hat{p}_t$ falls into a confidence interval. This standard 'belief band of inaction' is given by (6), where $z_{\alpha/2}$ is the appropriate quantile of a standard normal distribution.

$$p_0 - z_{\alpha/2}\sqrt{\frac{p_0(1-p_0)}{t}} < \hat{p}_t < p_0 + z_{\alpha/2}\sqrt{\frac{p_0(1-p_0)}{t}}. \tag{6}$$

We now demonstrate exactly the same $1/\sqrt{t}$ scaling effect from a very different perspective. For notational simplicity, we shall write $\hat{\sigma}_t^2 = \hat{p}_t(1-\hat{p}_t)$, and note that for large t, $\hat{\sigma}_t^2 \to p(1-p)$, in particular $\hat{\sigma}_t$ asymptotically does not depend on t.

Suppose our agent uses an estimator θ_t, based on the sample Y_i. She incurs two costs:

- A cost λ whenever θ_t changes

- A cost $\rho(\theta_t - \hat{p}_t)$, based on the error between θ_t and the MLE $\hat{p}_t$, paid at every time. We assume ρ is twice differentiable, convex and has a minimum $\rho(0) = 0$.

Remark 2. The cost ρ can be motivated in various ways. One approach is to treat the true probability p in a Bayesian fashion, and assume our agent faces a running cost $E[(p-\theta_t)^2|\mathcal{F}_t]$, that is, a cost depending on the distance of their estimate from the true (unknown) value. In this case, the MLE satisfies $\hat{p}_t = E[p|\mathcal{F}_t]$, and we can compute

$$E[(p-\theta_t)^2|\mathcal{F}_t] = E[(p-\hat{p}_t)^2|\mathcal{F}_t] + (\hat{p}_t - \theta_t)^2.$$

As the agent has no control over the term $E[(p-\hat{p}_t)^2|\mathcal{F}_t]$, the effective cost is given by $(\hat{p}_t - \theta_t)^2$, which is of the form considered.

From the theory of problems with transaction costs, see for example [2], the optimal policy is for the agent not to act until the error $W_t := \hat{p}_t - \theta_t$ leaves some interval. To a first approximation, which we consider more formally below, the interval is of the form $(-b, b)$, for some b to be determined. When W does leave this interval, the optimal strategy is to set $\theta_t = \hat{p}_t$, or equivalently $W_t = 0$(this is essentially because $\hat{p}_t$ is an unbiased estimate of p).

To find b, we first consider the behaviour of W between two sequential switching times $t_1 < t_2$. We can write

$$W_t = \hat{p}_t - \hat{p}_{t_1} = \frac{1}{t}\sum_{i=t_1+1}^{t}(Y_i - \hat{p}_{t_1}) \approx \frac{1}{t_1}\sum_{i=t_1+1}^{t}(Y_i - \hat{p}_{t_1}). \tag{7}$$

where the approximation is justified whenever $t_1^{-1} - t_2^{-1}$ is small. Hence W_t is approximately the sum of a sequence of mean-zero iid random variables, and so is well modelled as a random walk with up-probability p_{t_1}. Clearly, this approximation is better when t_1 is large and b is small (so switches occur more frequently).

Using the approximation of W as a random walk, we choose b to minimize expected costs. We have to trade off between our running cost and the cost of switching. For a time s, we try and evaluate the expected cost at time t, given our barrier strategy b. We first compute the running cost term.

Write $C_t(b) = \rho(\hat{p}_t - \theta_t) = \rho(W_t)$ when θ_t is determined using a boundary b. From our assumptions on u, provided W_t is not too large (which will happen whenever b is small or t is large), we can approximate with Taylor's theorem $\rho(W_t) \approx \gamma W_t^2$ for some constant γ.

Assuming our agent will be active over a long horizon, it is the long-run average value of this cost which is important. As b may change through time, it is natural to rescale our random walk, and see that the stationary distribution of W/b approximately has 'triangular' density

$$g(W/b) = \begin{cases} 1 + W/b & \text{if } -1 < W/b \leq 0 \\ 1 - W/b & \text{if } 0 < W/b \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

This can be seen by the facts that:

- $g(-1) = g(1) = 0$,
- The density integrates to unity.
- From considering the possible paths of W, except at $w = 0$, the only way for W to reach w is from being previously at either $w - \frac{1-\hat{p}_t}{bt_1}$ and observing $Y = 1$, or at $w + \frac{\hat{p}_{t_1}}{bt_1}$ and then observing $Y = 0$. In the stationary distribution, this implies

$$g(w) = \hat{p}_{t_1} g\Big(w - \frac{1-\hat{p}_{t_1}}{bt_1}\Big) + (1-\hat{p}_{t_1}) g\Big(w + \frac{\hat{p}_{t_1}}{bt_1}\Big).$$

  This equation is solved by linear functions g, and only by linear functions in the limit $bt_1 \to \infty$. Given the relationship does not need to hold at $w = 0$, we obtain a triangular density.

The density $g(.)$ has variance 1/6 and substituting,

$$E[C_t(b)|\mathcal{F}_s] \approx \gamma \frac{b^2}{6}$$

for $s \ll t$.

We now seek to understand the expected switching cost, $E[C_t^\lambda(b)|\mathcal{F}_s]$, where $C^\lambda(b) = \lambda$ if $|W_t| \geq b$ and zero otherwise. Using our rescaled random walk W/b, we need to find the probability of W/b hitting ± 1 at a time $t \gg s$. As W/b is approximately a random walk restarted at zero, this is approximately $1/E[\tau|\mathcal{F}_{t_1}]$, where $\tau = t_2 - t_1$ is the time taken to hit ± 1 from zero.

To calculate $E[\tau|\mathcal{F}_{t_1}]$, heuristically we calculate when the standard deviation of W/b equals 1. (Formally, this can be verified using the optional stopping theorem for the martingale part of $(W/b)^2$.) Starting at t_1, from (7), W_t/b has approximate variance

$$\mathrm{Var}(W_t/b) \approx \frac{1}{b^2 t_1^2} \sum_{i=t_1+1}^{t} \hat{\sigma}_{t_1}^2 = (t - t_1) \cdot \frac{\hat{\sigma}_{t_1}^2}{b^2 t_1^2}$$

so to find the average time to hit ± 1,

$$\sqrt{E[\tau|\mathcal{F}_{t_1}]}\Big(\frac{\hat{\sigma}_{t_1}}{bt_1}\Big) = 1 \quad \Rightarrow \quad E[\tau|\mathcal{F}_{t_1}] = \frac{(bt_1)^2}{\hat{\sigma}_{t_1}^2} \approx \frac{(bt_1)^2}{\hat{\sigma}_s^2},$$

provided s is sufficiently large that $\hat{\sigma}_{t_1}^2 \approx \hat{\sigma}_s^2$.

For $t \gg s$, with $t_1 < t \leq t_2$, the probability of hitting the barrier at t, incurring cost λ, is $1/E[\tau|\mathcal{F}_{t_1}]$, so

$$E[C_t^\lambda(b)|\mathcal{F}_s] \approx \lambda \Big(\frac{\hat{\sigma}_s}{bt}\Big)^2.$$

We can now minimize our expectations of long-run future costs $\Omega(b) = E[C_t^\lambda(b)|\mathcal{F}_s] + E[C_t(b)|\mathcal{F}_s]$ in a pointwise fashion:

$$\frac{d}{db}\Omega(b) = 0 \quad \Longleftrightarrow \quad \frac{d}{db}\left(\lambda\left(\frac{\hat{\sigma}_s}{bt}\right)^2 + \gamma\frac{b^2}{6}\right) = 0$$

which yields

$$b = \left(\frac{6\lambda}{\gamma}\right)^{1/4}\sqrt{\frac{\hat{\sigma}_s}{t}} = \chi\frac{\hat{\sigma}_s}{\sqrt{t}} \text{ where } \chi = \left(\frac{6\lambda}{\gamma}\right)^{1/4}\hat{\sigma}_s^{-1/2}.$$

Therefore, we obtain a bandwidth b for the belief band of inaction that is proportional to $1/\sqrt{t}$. As in the continuous time case, we also observe that the width χ of our band of inaction involves the square root of the data quality ratio $\propto \sigma_s^{-1/2}$. (To see that this corresponds to the data quality ratio, see that $\mathrm{Var}(\hat{p}_t - \hat{p}_{t-1}) \approx \hat{\sigma}_t/t$, while the variance of $\hat{p}_t$ is $\hat{\sigma}_t^2/t$. The ratio $\hat{\sigma}_t^{-1}$ then corresponds to the term H_t/R_t in the Kalman–Bucy dynamics.)

This demonstrates the equivalence with a hypothesis test in the case of a proportion, as desired. Using this b, we have

$$\Omega(b) = \frac{2\hat{\sigma}_t}{t}\sqrt{\frac{\lambda\gamma}{6}}$$

which suggests the $\lambda^{1/2}$ scaling of our asymptotic approximation in continuous time. We also note that $b \to 0$ and $bt \to \infty$ as $t \to \infty$, implying that our assumptions can be justified over long horizons.