# 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

**Abstract**

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

*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.

# 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 (sub-optimal) Maximum Likelihood Estimator (MLE) breaches it. We show that, to a first order approximation, this band of inaction is increasing with the inverse of the square root of the sample size, making it equivalent to a confidence interval and therefore to a two-sided hypothesis test. This is of interest in its own right, but the result widens the applicability of belief formation models based on hypothesis tests, such as Menzies and Zizzo (2009).

In section 2 we illustrate the key ideas informally using a discrete time estimator of a proportion. A proportion is an object of interest in many models, being an equilibrium outcome for recruitment processes, markets with informed and uninformed traders, dynamic goods' pricing models, and discrete choice models (Kirman 1993; Grossman and Stiglitz 1980; Brock and Hommes 1997; Brock and Durlauf 2001). In models of search and labour, the proportion of types is often stipulated in the model setup, as it was in the seminal assymetric-information 'market for lemons' paper (Akerlof 1970). Proportions are a way of representing probabilities in experiments (Khaw et al. 2017) and they are used by policymakers to influence social norms. Benabou and Tirole (2011) give the examples of estimating proportions like the incidence of tax compliance, electricity usage and recycling.

To set the scene, we consider estimating a parameter $p \in [0, 1]$, sampling from Bernoulli trials $X_i$ 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 $X_1, ..., X_t$.

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

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

A two-sided hypothesis test of size $\alpha$ 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 (1), where $z_{\alpha/2}$ is the appropriate quantile of a standard normal distribution.

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

# 2 Proof in Discrete Time

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 $\theta_t$, based on the sample $X_i$. She incurs two costs:

- A cost $\lambda$ whenever $\theta_t$ changes
- A cost $u\left(\frac{\theta_t - \hat{p}_t}{\hat{\sigma}_t}\right)$, based on the standardized error between $\theta_t$ and the MLE $\hat{p}_t$, paid at every time. We assume $u$ is twice differentiable and has a minimum $u(0) = 0$.

This type of problem is conceptually similar to problems with investment under fixed transaction costs, and we shall explore how the existing results can be translated into this context. We assume no discounting of costs.

*Remark* 1. The cost $u$ 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 $E[(p-\theta_t)^2|\mathcal{F}_t]/\hat{\sigma}_t^2$, that is, a cost depending on the distance of their estimate from the true (unknown) value, scaled against the best error. In this case, the MLE satisfies $\hat{p}_t = E[p|X_1, ..., X_t]$, 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 $(\hat{p}_t-\theta_t)^2/\hat{\sigma}_t^2$, which is of the form considered.

From the theory of problems with transaction costs, 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 $W_{t_2} = 0$ (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

$$W_t = \hat{p}_t - \hat{p}_{t_1} = \frac{1}{t}\sum_{i=t_1+1}^{t}(X_i - \hat{p}_{t_i}) \approx \frac{1}{t_1}\sum_{i=t_1+1}^{t}(X_i - \hat{p}_{t_i}). \tag{2}$$

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.

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 $C_t(b) = u(\hat{p}_t - \theta_t) = u(W_t)$ when $\theta_t$ 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 $u(W_t) \approx \gamma W_t^2/\hat{\sigma}_t^2$ for some constant $\gamma$.

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 $W/b$ 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, and the density obeys

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

from considering the possible paths of $W$. Assuming $bt_1$ is large, this yields the stated triangular density. The density $g(.)$ has variance $1/6$ and substituting,

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

for $s \ll t$. Provided $s$ is large, the difference between $\hat{\sigma}_s^2$ and $\hat{\sigma}_{t_1}$ will be small, justifying the final approximation.

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 $\pm 1$ at a time $t \gg s$. As $W/b$ is approximately a random walk restarted at zero, this is approximately $1/E[\tau]$, where $\tau = t_2 - t_1$ is the time taken to hit $\pm 1$ 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 (2), $W_t/b$ has approximate variance

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

so to find the average time to hit $\pm 1$,

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

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

For $t \gg s$, the probability of hitting the barrier, incurring cost $\lambda$, is $1/E[\tau|\mathcal{F}_s]$, so

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

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

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

which yields

$$b = \chi \frac{\hat{\sigma}_s}{\sqrt{t}}, \text{ where } \chi = (6\lambda/\gamma)^{1/4}. \tag{3}$$

Therefore, we obtain a bandwidth $b$ for the belief band of inaction that is proportional to $1/\sqrt{t}$. This demonstrates the equivalence with a hypothesis test in the case of a proportion, as desired. We also note that $bt \to \infty$ as $t \to \infty$, implying that our assumptions that $t$ and $bt$ are large can be justified over long horizons.

# References


Akerlof, G. (1970), 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism', *Quarterly Journal of Economics*, 84(3), August, 488-500.

Benabou R. and J. Tirole (2011), 'Laws and Norms', *NBER working paper,* 17579, November.

Brock, W. and S. Durlauf (2001), 'Discrete Choice with Social Interactions', *The Review of Economic Studies*, 68(2), 235-260.

Brock, W. and C. Hommes (1997) 'A Rational Route to Randomness', *Econometrica*, 65(5), 1059–1095

Grossman S. and J. Stiglitz (1980), 'On the Impossibility of Informationally Efficient Markets' *The American Economic Review*, 70(3), June, 393-408.

Khaw, M., Stevens, L. and M. Woodford (2017), 'Discrete Adjustment to a Changing Environment: Experimental Evidence' *Journal of Monetary Economics*, 91, 88-103.

Kirman, A. (1993), 'Ants, Rationality, and Recruitment', *The Quarterly Journal of Economics*, 108(1),137–156.

Menzies, G. and D. Zizzo (2009), 'Inferential Expectations', *B. E. Journal of Macroeconomics*, 9, 1-25.