Brownian Motion and Normally Distributed Beliefs*

Godfrey Keller†

June 2006, revised May 2011

Abstract

We consider a model that has the following features: (a) time is continuous; (b) agents are observing Brownian motion with an unknown drift; (c) their prior belief about the drift is that it is normally distributed with a given mean and variance. (In almost all other economic models with (a) and (b), the prior belief has a discrete, two-point distribution.) As a consequence of (b) and (c), the posterior belief is also characterised by a normal distribution, and we derive the incremental mean and variance from the prior and the observations. Further, we prove that the belief converges to the true value almost surely.

Keywords: Bayesian Learning, Controlled Diffusion.

JEL Classification Numbers: D83, O33.

---

*I would like to thank Boyan Jovanovic, Sven Rady, and John Thanassoulis for a useful dialogue and helpful discussions.

†Department of Economics, University of Oxford, Manor Road Building, Oxford OX1 3UQ, UK. e-mail: godfrey.keller@economics.ox.ac.uk

1 Introduction

There are several notable articles in the economics literature that consider the problem of agents observing Brownian motion with known volatility but unknown drift – be it the value of a job match to a firm and worker, the perceived quality of a new consumption good, the level of demand for a firm's product in a market – and the overwhelming majority of these assume that the drift is just one of two values – *High* or *Low*; see, for example, Felli & Harris (1996), Bergemann & Välimäki (1997), Keller & Rady (1999), Bolton & Harris (1999), and Moscarini (2005).[1] In these models, the distribution function of beliefs characterising the uncertainty has a discrete, two-point support, and so is captured by a single state variable, namely the level of optimism that the drift is high. How this belief evolves is given in Liptser & Shiryayev (1977, Theorem 9.1), and convergence properties can be found in Karlin & Taylor (1981, Chapter 15, Sections 6 & 7).

In other models, it might not be appropriate for there to be just two possible values of the unknown drift (or indeed, any finite or countable number). Here, we make the same assumptions as in the job-matching model in Jovanovic (1979), namely, that the drift was drawn by nature from a continuous distribution that is normal with known mean and variance, so the prior belief of the agents is normally distributed. It follows from Chernoff (1968) that the agents' posterior belief is also normally distributed, with an updated mean and variance.[2] In the main section, we derive the incremental changes to the mean and variance, and provide two alternative representations of the incremental mean – see equations (4) and (5); moreover, we establish the convergence properties of the agents' belief.[3]

2 The Model

At each moment in time, the agents observe $dX(t) = \mu\, dt + \sigma\, dz(t)$ where μ is fixed but unknown, σ is fixed and known, and dz is the increment of a Wiener process. At time

[1] Felli & Harris (1996) study wage determination in the presence of firm-specific human capital, based on a job-matching model of the labour market, and Moscarini (2005) embeds a microeconomic job-matching model in a macroeconomic equilibrium search environment; Bergemann & Välimäki (1997) analyse the diffusion of a new product of uncertain value in a duopolistic market; Keller & Rady (1999) study experimentation by a monopolist facing an unknown demand curve which is subject to random changes. Bolton & Harris (1999) extend the two-armed bandit problem to a many-agent setting where each agent can now learn from the experimentation of others – theirs is rather an abstract theoretical model, as is the one studied here.

[2] In discrete time, it is well known that the normal distribution is a conjugate family when observations are made with normally distributed noise – see, for example, DeGroot (1970, Chapter 9).

[3] Jovanovic (1979) does not have the correct equation for updating the mean (equation 9 on p.979), nor is convergence demonstrated.

$t = 0$ the agents have a prior belief that $\mu \sim \mathrm{N}(m_0, s_0)$, and update this belief in a Bayesian fashion in light of the observations. Following Chernoff (1968, Lemma 4.1), at time t they have a posterior belief that $\mu \sim \mathrm{N}(m(t), s(t))$, with $m(0) = m_0$ and $s(0) = s_0$:

$$m(t) = s(t)\left[m_0 s_0^{-1} + \sigma^{-2}\int_0^t dX(u)\right], \tag{1}$$

$$s(t) = \left(s_0^{-1} + \sigma^{-2} t\right)^{-1}. \tag{2}$$

The variance of the posterior belief, $s(t)$, is deterministic and decreases towards 0; the mean of the posterior belief, $m(t)$, however, is a random variable, and, as we show in section 2.1, it is a martingale *given the information available to the agents.* In section 2.2, we show that $m(t)$ is itself normally distributed, with a mean that tends to the true value of the unknown drift and a variance that tends to 0.

2.1 Evolution of Beliefs

Equation (1) implies that $dm(t) = m(t)s(t)^{-1}\,ds(t) + s(t)\sigma^{-2}\,dX(t)$, and equation (2) implies that $ds(t) = -s(t)^2\sigma^{-2}\,dt$, so

$$\begin{aligned} dm(t) &= -m(t)s(t)\sigma^{-2}\,dt + s(t)\sigma^{-2}\,dX(t) \\ &= s(t)\sigma^{-2}\left(dX(t) - m(t)\,dt\right). \end{aligned} \tag{3}$$

These are the incremental mean and variance[4] to be used when applying Itô's lemma, for example, although one of the alternative representations of $dm(t)$ given in equations (4) and (5) is usually more convenient.

Replacing $dX(t)$ in equation (3) gives us the representation

$$dm(t) = s(t)\sigma^{-2}\left([\mu - m(t)]\,dt + \sigma\,dz(t)\right), \tag{4}$$

and the term in square brackets in equation (4) shows that, as long as there is something that can be learned (i.e. $s(t) > 0$), $m(t)$ is pulled towards the true value, μ. Even though $s(t) \downarrow 0$, this does not imply that $m(t)$ converges, since the variance of Brownian motion increases linearly with t; as we shall see, we need the stronger condition that $s(t)^2 t \downarrow 0$.

Note that equation (3) or (4) can be written as

$$dm(t) = s(t)\sigma^{-1}\,d\bar{z}(t), \tag{5}$$

where

$$d\bar{z}(t) = \sigma^{-1}\left(dX(t) - m(t)\,dt\right) = \sigma^{-1}\left([\mu - m(t)]\,dt + \sigma\,dz(t)\right) \tag{6}$$

[4]Cf. Liptser & Shiryayev (1977, Theorem 10.1).

defines a Wiener process (conditional on the information available to the agents) that is related to that in the data generating process, and called the *innovation* process.[5] The representation in equation (5) shows that it is changes in $\bar{z}(t)$ that lead to revisions of $m(t)$, and equation (6) shows that any such change depends on the difference between the agents' observation and their expectation. Also, note the difference in perspective between the two representations of $dm(t)$: equation (5) is from that of the agents, and should be used when modelling their decisions; equation (4) is from that of 'nature' who knows the unobserved parameter μ, and should be used by the modeller when considering convergence of beliefs, or when simulating sample paths of the evolving mean for example.

2.2 Convergence of Beliefs

To find the distribution of $m(t)$, one could solve the stochastic differential equation (4); however, as we already have equation (1), we can replace $dX(u)$ there and integrate:

$$\begin{aligned} m(t) &= s(t)\left[m_0 s_0^{-1} + \sigma^{-2}\int_0^t \left(\mu\,du + \sigma\,dz(u)\right)\right] \\ &= s(t)\left[m_0 s_0^{-1} + \sigma^{-2}\mu t + \sigma^{-1}z(t)\right] \end{aligned} \tag{7}$$

showing that the random variable $m(t)$ is such that

$$m(t) \sim \mathrm{N}\left(\frac{m_0 s_0^{-1} + \mu\sigma^{-2}t}{s_0^{-1} + \sigma^{-2}t}, \frac{\sigma^{-2}t}{\left(s_0^{-1} + \sigma^{-2}t\right)^2}\right) \quad \rightarrow \quad \mathrm{N}(\mu, 0) \text{ as } t \rightarrow \infty,$$

i.e. $m(t)$ converges to a degenerate random variable, and, since $s(t)$ converges to 0, the state vector $\langle m(t), s(t)\rangle \rightarrow \langle \mu, 0\rangle$ as $t \rightarrow \infty$. This shows that each sample path $\{m(t)\}_{t\geq 0}$ converges to μ with probability 1.

Even the agents in the model know that they will learn the true value of μ eventually – it's just that they don't know μ now. As a consequence, they use the *un*conditional distribution of $m(t)$, namely $m(t) \sim \mathrm{N}\left(m_0, s_0 - s(t)\right)$, which can be calculated either from equation (7) with $\mu \sim \mathrm{N}(m_0, s_0)$ at $t = 0$, or by solving the stochastic differential equation (5), ignoring the fact that $\bar{z}$ and X are related.[6]

[5] This terminology is used in Liptser & Shiryayev (1977), where they show that the processes $\bar{z}$ and X are informationally equivalent.

[6] From equation (5), $m(t)$ is normally distributed with $\mathrm{E}[dm(t)] = 0$, so $\mathrm{E}[m(t)] = m(0) = m_0$; and $\mathrm{Var}[dm(t)] = s(t)^2\sigma^{-2}\,dt$, so $\mathrm{Var}[m(t)] = \int_0^t s(u)^2\sigma^{-2}\,du = -\int_0^t ds(u) = s(0) - s(t) = s_0 - s(t)$.

References

Bergemann, D. and Välimäki, J. (1997): "Market Diffusion with Two-Sided Learning," *RAND Journal of Economics*, **28**, 773–795.

Bolton, P. and Harris, C. (1999): "Strategic Experimentation," *Econometrica*, **67**, 349–374.

Chernoff, H. (1968): "Optimal Stochastic Control," *Sankhyā*, **30**, 221–252.

DeGroot, M. (1970): *Optimal Statistical Decisions*. New York: McGraw Hill.

Felli, L. and Harris, C. (1996): "Learning, Wage Dynamics and Firm-Specific Human Capital," *Journal of Political Economy*, **104**, 838–868.

Jovanovic, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, **87**, 972–990.

Karlin, S. and Taylor, H.M. (1981): *A Second Course in Stochastic Processes*. New York: Academic Press.

Keller, G. and Rady, S. (1999): "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, **66**, 475–507.

Liptser, R.S. and Shiryayev, A.N. (1977): *Statistics of Random Processes I*. New York: Springer-Verlag.

Moscarini, G. (2005): "Job Matching and the Wage Distribution," *Econometrica*, **73**, 481–516.