Optimal Sequential Experimentation

Jorge D. Ramos-Mercado*

November 22, 2023

Abstract

I study flexible, sequential experimentation. A decision-maker ($\mathcal{DM}$) learns about a state by managing a jump-diffusion process with state-dependent dynamics. He controls the diffusion's precision and the state-dependent arrival rate of jumps, but faces convex flow costs increasing in the amount of acquired information. I find that an optimal pure-diffusion experiment exists. Intuitively, the $\mathcal{DM}$ cannot separately manage the precision and composition of acquired information. Hence, any optimal experiment (with jumps) equalizes the marginal benefits and costs of acquiring information from both types of processes. In contrast, if the $\mathcal{DM}$ did not have to explicitly experiment, Zhong (2021) finds it optimal to learn from learning from a pure-jump.

I study a flexible sequential experimentation problem. A decision-maker ($\mathcal{DM}$) learns about a payoff-relevant state by managing a jump-diffusion signal process. He controls the diffusion's precision and state-dependent arrival rate of jumps. The diffusion describes imprecise but frequently arriving (i.e., Gaussian) information; meanwhile, jumps model Poisson information: precise but infrequently arriving. If the $\mathcal{DM}$ is impatient and faces flow costs that convexly increase with the amount of generate information, how should he experiment? I find restricting the $\mathcal{DM}$ to only acquiring Gaussian information is without loss of generality.

Wald (1947) first studied sequential experimentation. He studied until when a $\mathcal{DM}$ should acquire noisy signals (i.e., experiments) about a relevant state prior to making an

*Ramos-Mercado: EGADE Business School, Tecnológico de Monterrey, ramos237@tec.mx.

irreversible decision. Moscharini and Smith (2001), *MS* from henceforth, extends this setting by allowing the $\mathcal{DM}$ to control the precision of a pure-diffusion signal process where only the drift depends on the state. *MS* finds that precision increases in the expected payoff from making a decision.

This generalization, nevertheless, imposes that the $\mathcal{DM}$ must learn gradually. This is a restriction, because an experiment (seen as a Lévy signal process) may generate independent and infrequent jumps whose arrival rate depends on the state. In other words, a more general set of experiments also allows the $\mathcal{DM}$ to learn discontinuously. Would expanding the space of experiments make the $\mathcal{DM}$ better off? Zhong (2022), *Z22* from henceforth, finds that it does.

Rather than extending *MS* , *Z22* considers a reduced-form approach. He assumes that the $\mathcal{DM}$ directly picks a Martingale, Lévy process for his beliefs about the state. Meanwhile, the cost of experimenting increases in the amount of information acquired—itself, a function of beliefs. This manner of modeling experimentation costs follows from the rational inattention literature e.g., Sims (2003), Hébert and Woodford (2021), Caplin et al (2022), Macowiak et al (2023).[1] *Z22* finds a unique, robust prediction: it is optimal to learn from a pure-jump process seeking to confirm the most likeky state.

In order to replicate such learning process, the $\mathcal{DM}$ must run an experiment producing noise-less data. In the real-world, however, said experiments do not exist. Moreover, managing an experiment's data precision (i.e., data cleaning) is a central concern. Data cleaning, for example, takes up roughly 60 percent of data scientists (Press 2016) and 80 percent of clinical researchers (Rozario et al 2017) time. In contrast, I extend *MS* 's model by allowing for experiments generating jumps, but assume costs are as in *Z22* . This allows me to model costs in a consistent fashion across experiments.

I find that a $\mathcal{DM}$ forced to explicitly experiment cannot be made strictly worse off by disallowing him from acquiring Poisson information i.e., running an experiment with jumps. The intuition goes as follows. The $\mathcal{DM}$ completes two, inseparable tasks: manage data precision and decide how much to learn from the diffusion and from jumps. As a consequence, if an optimal experiment generates jumps, the $\mathcal{DM}$ equalizes the marginal

[1] Said approach has garnered critique from Denti et al (2022) since this cost structure is inconsistent with a model of experiment specific costs.

benefit and costs of acquiring both types of information. Suppose that an optimal experiment produces jumps. I exploit the fact that experiments run in continuous-time and the information based cost structure to find a payoff-equivalent experiment that never generates jumps.

The rest of the paper proceeds as follows. Section 1 presents the model. Section 2 then states the result. Section 3 concludes.

1 Model

I now present the model. A Bayesian decision-maker ($\mathcal{DM}$) with Bernoulli preferences picks from a finite set of alternatives $A, |A| = 2, 3, \ldots$. His payoff also depends on an unknown state $x \in X \equiv \{x_i\}_{i=1}^n$ for $n = 2, 3, \ldots$: payoffs are $u : A \times X \to \mathrm{R} \gg 0$. Meanwhile, his initial beliefs are $p \equiv (p_i)_{i=1}^n \in \Delta^{n-1}$ where $\forall i, p_i \equiv \Pr(x = x_i) > 0$.

Next, the $\mathcal{DM}$ is note required to decide from the outset. Instead, he can experiment and make a decision at some time $T \in [0, \infty)$. The details of how he experiments are presented below. However, if at time T he holds beliefs $p_T \in \Delta^{n-1}$, his payoff from making a decision are

$$F(p_T) \equiv \max_{a \in A} \sum_{i=1}^{n} p_{iT} u(a, x_i). \qquad \text{(Terminal Payoffs)}$$

I assume that for each pair of states x_i, x_j, $\mathrm{argmax}_{a \in A}\, u(a, x_i) \cap \mathrm{argmax}_{a \in A}\, u(a, x_i) \neq \emptyset$ if and only if (iff) $x_i = x_j$. This ensures that learning about the state is payoff relevant.

1.1 Information Acquisition Problem

Signals I now describe the experimentation problem. The $\mathcal{DM}$ picks a continuous-time signal process and a signal-adapted stopping time $T < \infty$. The stopping time denotes when experimentation stops and the $\mathcal{DM}$ makes an irreversible decision. An admissible signal (i.e., experiment) $s \equiv (s_t)$ is a jump-diffusion process such that $s_0 = 0$ and at each time $t \in [0, \infty)$

$$\mathrm{d}s_t = \mu(x)\,\mathrm{d}t + \frac{\mathrm{d}B_t}{\sqrt{h_t}} + \mathrm{d}N_t \tag{1}$$

The process $B \equiv (B_t)$ is a Brownian motion with precision $h_t \gg 0$ and drift $\mu(x)$ where $\mu : X \to (-\infty, \infty)$ is an injective function. Meanwhile, $N \equiv (N_t)$ is a compensated jump process that is independent of B and jumps by 1 at a rate of $\lambda_{it} \geq 0$ at time t iff $x = x_i$. Lastly, I assume that $(h_t, (\lambda_{it})_{i=1}^n)$ satisfy the standard the standard Lipschitz condition which ensuring that (s_t) admits a weak solution[2], but no additional restrictions are made on the set of feasible parameters.[3]

Information and Costs I now model an experiment's flow costs. The costs of running an experiment increase in the flow amount on information generated. To do so, I first derive a measurement of the flow amount of information generated as in *Z22* .

Let $H : \Delta^{n-1} \to \mathrm{R} \in C^2$ (e.g., entropy) be a strictly concave function and consider the belief process $(p_t \equiv p_{it}(x = x_i \mid \{s_\tau : \tau \in [0,t]\})) \subset [0,1]$. The flow amount of information generated by (s_t) at time t is $I_t \equiv -\mathcal{L}H(p_t)$ where $\mathcal{L}(\cdot)$ is the infinitesimal generator for (p_t) i.e., for each function f, $\mathcal{L}f(p_t) \equiv \lim_{\mathrm{dt}\to 0} \frac{f(p_t) - f(p_{t-\mathrm{dt}})}{\mathrm{dt}}$ if said limit exists. Lastly, the flow cost of experimenting at time t is $c(I_t)$ for some $c(\cdot)$ being a strictly increasing, convex, and twice differentiable function. Lastly, I assume that $(I_t) \subset [0, \bar{I}]$ for some constant $\bar{I} > 0$. This means that the $\mathcal{DM}$ faces a constraint on how much from information he can acquire.

Payoffs I now describe payoffs. If the $\mathcal{DM}$ picks signal s (that generates processes (p_t) and (I_t)) and an s-adapted stopping time T , then at time t ($\leq T$) expected payoffs are

$$V_t(s,T) \equiv E_t\bigg[e^{-r(T-t)}F(p_T) - \int_t^T c(I_t)e^{-(\tau-t)}\mathrm{d}\tau \mid \{s_\tau : \tau \in [0,t]\}\bigg] \tag{Payoffs}$$

[2]See for example Oksendal and Sulem (2019) among others

[3]I choose this model for its parsimony. One could allow the decision-maker to observe a multidimensional Jump-diffusion and for jumps to take on a finite number of jumps—for instance. However, the results would extend, but the statements in the results would be more complex.

Hence, the $\mathcal{DM}$'s problem given initial beliefs p is

$$V(p) = \max_{s,T} V_0(s,T). \qquad \text{(Unconstrained Problem)}$$

Alternatively, the $\mathcal{DM}$ may be forced to only acquire pure-diffusion experiments:

$$U(p) = \max_{s,T} V_0(s,T) \text{s.t. } \forall i, t, \lambda_{it} = 0 \text{a.s.} \qquad \text{(Constrained Problem)}$$

I conclude by noting that $U(p) \leq V(p)$ since all signals feasible in the restricted problem are also feasible in the unrestricted problem.

2 Results.

I now present my results. The proof has four parts. First, I explicitly derive beliefs. Next, I use standard techniques from stochastic calculus to derive a value function for the restricted and unrestricted problems. I then conduct a sequence of change of variables which makes my main result straightforward.1

2.1 Belief dynamics

I first characterize how a function of Bayes beliefs, derived from (s_t), changes over time.

Lemma 2.1. *Fix $s = (s_t)$. Let (p_t) be the s-adapted Bayes consistent beliefs. Then for each $f : \Delta^{n-1} \to \mathrm{R} \in C^2$, then*

$$\begin{aligned}\mathcal{L}f(p_t) &= h_t \sum_{ij} \frac{p_{it}p_{jt}}{2}(\mu_i - \mu_t)(\mu_j - \mu_t) f_{ij}(p_t) + \lambda_t[f(\nu_t) - f(p_t) - \nabla f(p_t) \cdot (\nu_t - p_t)] \\ &= \frac{h_t}{2} tr[\tilde{\mu}_t' Hf(p_t) \tilde{\mu}_t] + \lambda_t[f(\nu_t) - f(p_t) - \nabla f(p_t) \cdot (\nu_t - p_t)] \quad (2)\end{aligned}$$

where $f_{ij}(p_t) = \partial_{p_i}\partial_{p_j} f(p_t)$, $\lambda_t \equiv \sum_i p_{it}\lambda_{it}$, $\mu_t \equiv \sum_i p_{it}\mu_i$, $\nu_t \equiv (p_{it^-}\frac{\lambda_{it-}}{\lambda_{t^-}})_{i=1}^n$, $\tilde{\mu}_t \equiv (p_{it}(\mu_i - \mu_t))_{i=1}^n$, $p_{t^-} = \lim_{dt\to 0} p_{t-dt}$, $f(\cdot)$'s hessian at p_t is $Hf(p_t)$, and $tr(\cdot)$ is a trace.

I now sketch the proof, but delegate the derivation to appendix A.1. Note that the jump and diffusion process are independent. This allows me to approximate each as separate binomials. For some small, time interval dt> 0, the diffusion jumps up by $\pm\sqrt{\mathrm{dt}/h_t}$ with probability $\approx (1 \pm \mu_i\sqrt{h_t\mathrm{dt}})/2$ when $x = x_i$. Note that I denote $\mu_i \equiv \mu(x_i)$ for each i. Likewise, the jump process jumps by 1 with probability $\approx \lambda_{it}$dt, but it otherwise remains equal to 0. I then approximate the change in a function of beliefs $f(\cdot) \in C^2$ in two parts: there is or there is no jump. When there is a jump, Bayes rule updated beliefs solely on the relative arrival rate of jumps; meanwhile, there is no jump, the change in f is approximated via Taylor's rule.

Next, I make two observations. First, if at some time $t \geq 0$, $\forall x_i, \lambda_{x_it} = 0$, then I can define $\nu_t = p_{t^-}$ since jumps are a 0-probability event. The second observation is that the dynamics when the $\mathcal{DM}$ cannot elicit jumps follow immediately and are stated below.

Corllary 2.2. *Let $s = (s_t)$ be a signal such that at each time t and $x = 0, 1$, $\lambda_{xt} = 0$ almost surely (a.s.), then for each twice continuously differentiable function $f(\cdot)$, it holds that $\mathcal{L}f(p_t) = \frac{h_t}{2}tr[\tilde{\mu}_t' Hf(p_t)\tilde{\mu}_t]$.*

2.2 Re-formulating the Experimentation Problems

Now that the beliefs dynamics are well-defined, I define the cost function, derive the Hamilton-Jacobi-Bellman (HJB) that the decision-maker's problem must solve, and then reformulate it in a more useful fashion. First, I characterize the costs function. Fix some signal $s = (s_t)$ process, then at each time t the amount of information generate is $I_t \equiv \frac{h_t}{2}\mathrm{tr}[\tilde{\mu}_t' H[H(p_t)]\tilde{\mu}_t] + \lambda_t[H(p_t) - H(\nu_t) - \nabla H(p_t) \cdot (p_t - \nu_t)]$. Now that I derived an expression for the amount of information generated, the flow costs are $c(\mathrm{I}_t)$. The formula above further illustrates several points of note. First, the signal precision enters linearly into the total amount of information and separable from the information derived from jumps. This is a feature of the continuous-time modeling choice and plays a key role in the result below.

Next, I derive an expression for the HJB describing the $\mathcal{DM}$'s optimal experimentation problem. The $\mathcal{DM}$'s problem can be written as a function of his beliefs p_t at each time t. By the principle of optimality, if at some belief $p_t \in \Delta^{n-1}, V(p_t) > F(p_t)$, then $V(\cdot)$

satisfies

$$rV(p_t) = \max_{\phi_t} -\frac{h_t}{2}\text{tr}[\tilde{\mu}_t' HV(p_t)\tilde{\mu}_t] + \lambda_t[V(\nu_t) - V(p_t) - \nabla V(p_t)\cdot(\nu_t - p_t)] - c[\text{I}(\phi_t, p_t)] \quad (3)$$

s.t. $\phi_t \equiv (h_t, (\lambda_{it}))\ h_t > 0, \forall i, \lambda_{it} \geq 0, \text{I}(\phi_t, p_t) \leq \bar{I}$.

Oksendal and Sulem (2019) establish that the HJB equation, at least, admits a viscosity solution. This is because the $\mathcal{DM}$'s problem reduces to picking a locally Lipschitz collection of parameter process. They even extend the existence proof to a much broader set of problems than the one studied in this paper.

This HJB equation is far too general to make any useful insights. Instead, I consider a change of variables that clarifies the structure of the value function. Assume that the $\mathcal{DM}$ picks β_t, ν_t, and j_t such that $j_t \equiv \lambda_t[H(p_t) - H(\nu_t) - \nabla H(p_t)\cdot(p_t - \nu_t)]$ and $\beta_t \equiv \frac{h_t}{2}\text{tr}[\tilde{\mu}_t' H[H(p_t)]\tilde{\mu}_t]$ Then the flow of acquiring information is $c(I_t) = c(\beta_t + j_t)$. The generator of any function $f : \Delta^{n-1} \to \text{R} \in C^2$ becomes

$$\mathcal{L}f(p_t) = \beta_t L(f, p_t) + j_t G(f, p_t, \nu_t)$$

where $G(f, p_t, \nu_t) \equiv -\frac{f(\nu_t) - f(p_t) - \nabla f(p_t)\cdot(\nu_t - p_t)}{H(\nu_t) - H(p_t) + \nabla H(p_t)\cdot(\nu_t - p_t)}$ and $L(f, p_t) \equiv \frac{\text{tr}[\tilde{\mu}_t' H f(p_t)\tilde{\mu}_t]}{-\text{tr}[\tilde{\mu}_t' H[H(p_t)]\tilde{\mu}_t]}$. I can now state a more useful reformulation of the value function if $F(p_t) < V(p_t)$, then

$$rV(p_t) = \max_{j_t, \nu_t, \beta_t} \beta_t L(V, p_t) + j_t G(V, p_t, \nu_t) - c(\beta_t + j_t) \quad (4)$$

such that $j_t \geq 0, \nu_t \in \Delta^{n-1}, \beta_t > 0, j_t + \beta_t \leq \bar{I}$.

This formulation is useful since the problem separates the $\mathcal{DM}$'s problem into one of picking how much information to acquire from the diffusion (i.e., β_t), how much information to acquire from the jumps (j_t), and the posterior belief conditional on the arrival of a jump (ν_t).

In a similar fashion, the restricted problem forces that for each $i = 1, 2, \ldots, n\ \lambda_{it} = 0$ for certain. This implies that $j_t = 0$ and $\nu_t = p_t$. As a consequence, I can make the same derivations for the restricted problem and the value function $U(\cdot)$ satisfies that if $F(p_t) < U(p_t)$, then

$$rU(p_t) = \max_{\beta_t \in (0, \bar{I}]} \beta_t L(U, p_t) - c(\beta_t) \quad (5)$$

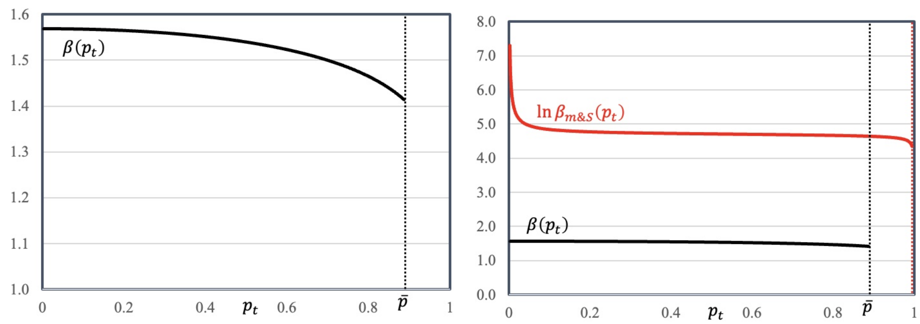

(a) ***Policy function*** *(b)* ***Comparison to* MS.**

Figure 1: ***Policy function and comparison to Moscarini and Smith (2001).***

I now characterize the value function $U(\cdot)$ and its optimal control in the lemma below.

Lemma 2.3. *The HJB stated in equation 5 has a unique solution $U(\cdot)$ that is twice continuously differentiable and admits a unique, optimal control that is a Markov function of p_t i.e., $\bar{\sigma} : \Delta^{n-1} \to \mathrm{R}_+$.*

The lemma is an immediate application of Theorem 1 and 2 from Strulovici and Szydlowski (2015). Now that the bound $\bar{I} < \infty$ simply ensures that the set of choices is a non-empty, compact space. Next, I illustrate the resulting policy function $\beta(\cdot)$ when $\bar{I}$ is large, $c(I) = I^2/2$, $x = 0, 1$, $p \equiv \Pr(x = 1) \in (0, 1)$, H is entropy, $a = 0, 1$, and terminal payoffs are $u(a, x) = a(x - 1/2)$. Said policy function is illustrated in the left-hand panel of figure 1. I find that the total amount of information acquired falls as beliefs approach a cutoff of $\bar{p} < 1$. When $p_t = \bar{p}$, the $\mathcal{DM}$ stops experimenting and picks $a = 1$.

Next, I consider an analogous version of the model in **MS** where the flow cost of managing a pure-diffusion signal with precision process (h_t) has flow costs $h_t^2/2$. I then estimate the flow amount of information acquired as measured by $H(\cdot)$ and take its log—since it makes the resulting value smaller. The right-hand panel in figure 1 illustrates (in red) the resulting log-amount of information acquired. I find that the $\mathcal{DM}$ in **MS**'s model would acquire more information and would only stop experimenting when his beliefs reach a higher cutoff belief—near 1.

2.3 Main Result

This section presents my main result. I will first state the result and then provide the proof.

Theorem 2.4. *Only generating Gaussian information is without loss:* $\forall p_t, U(p_t) = V(p_t)$.

This theorem implies that restricting the $\mathcal{DM}$ to picking a diffusion signal is without loss of generality. The proof goes as follows. First, it is immediate that for each belief $p_t \in \Delta^{n-1}$, it must be that $U(p_t) \leq V(p_t)$. This is because a control $\phi_U \equiv (\beta, I, \nu) : \Delta^{n-1} \to \mathrm{R}_{++} \times \mathrm{R}_{+} \times \Delta^{n-1}$ such that $h(\cdot)$ is the same control maximizing the constrained problem, $I(p_t) = 0$ and $\nu(p_t) = p_t$ is an admissible control for the general problem. Hence, the payoff from using said control when the current belief is p_t can be defined as $V(p_t; \phi_U)$ and it holds that $V(p_t, \phi_U) = U(p_t)$ and $V(p_t, \phi_U) \leq V(p_t)$.

What I need to show is that the inequality also holds in the opposite directions i.e., $U(p_t) \geq V(p_t)$. First observe that if at belief $p_t \in \Delta^{n-1}$, it holds that $V(p_t) = F(p_t)$, then $U(p_t) = F(p_t)$ since

$$\underbrace{F(p_t) \leq U(p_t)}_{\text{restricted } \mathcal{DM} \text{ can stop}} \leq \underbrace{V(p_t) = F(p_t)}_{\text{Unrestricted } \mathcal{DM} \text{ wants to stop}} .$$

This implies that it is without loss of generality to focus on the set of beliefs p_t for which the $\mathcal{DM}$ prefers to experiment when he is not restricted: $C \equiv \{p_t \in \Delta^{n-1} : V(p_t) > U(p_t)\}$. Next suppose that $\phi' = (\beta', I', \nu')\Delta^{n-1} \to \mathrm{R}_{++} \times \mathrm{R}_{+} \times \Delta^{n-1}$ attains the maximum of the HJB equation 4. At each belief $p_t \in C$, $\phi'(p_t)$ must satisfy two conditions—that are interior for large enough P. First, $\phi'(p_t)$ must satisfy the first order conditions. This implies that $\beta_t \gg 0$ for $p_t \in C$ and satisfies $c'(\beta_t + j_t) = L(p_t)$. Likewise, if $I'(p_t) > 0$, then it must satisfy that

$$\underbrace{c'(\beta_t + I_t)}_{\text{Marg. cost of Info.}} = \underbrace{G(V, p_t, \nu_t)}_{\text{Marg. benefit of breakthroughs}} = \underbrace{L(V, p_t)}_{\text{Marg. benefit of noisy data}}$$

Secondly, it must satisfy the principle of optimality i.e., for each $p_t \in C$, it holds that

$$\begin{aligned} rV(p_t) &= \beta'(p_t)L(V, p_t) + G(V, p_t, \nu'(p_t))I'(p_t) - c[\beta_t(p_t) + I_t(p_t)] \\ &= [\beta'(p_t) + I'(p_t)]L(V, p_t) - c[\beta_t(p_t) + I_t(p_t)] \end{aligned}$$

Notice that the second line follows from the observation that the marginal benefits both types of information must be equalized whenever the $\mathcal{DM}$ pick a strictly positive amount of both types of information. On the other hand, if $I'(p_t) = 0$, then the expression is unaffected by the $I'(\cdot)$ part of the expression. Alternatively, let $\bar{\phi} \equiv (\bar{\beta}, \bar{I}, \bar{\nu}) : [0,1] \to \mathrm{R}^2_+ \times [0,1]$ be defined for each p_t as $\bar{\beta}(p_t) = \beta'(p_t) + I'(p_t), \bar{I}(p_t) = 0$, and $\nu(p_t) = p_t$. Then, by construction, $\bar{\sigma}(p_t)$ satisfies the first order conditions for the optimization problem and for each $p_t \in C$, it holds that $V(p_t) = \bar{\beta}(p_t)L(V, p_t) - c(\bar{\beta}_t(p_t))$. Consequently, it holds that for each p_t $V(p_t; \bar{\phi}) = V(p_t)$. Likewise, $\bar{\phi}$ is admissible in the constrained problem. This implies that for each p_t, it holds that $V(p_t) = V(p_t; \phi) \leq U(p_t)$. This establishes the result.

3 Discussion

This paper studies optimal, flexible, sequential experimentation. I find that restricting the decision-maker to only run pure-diffusion experiments (i.e., learn gradually over time) is without loss of generality. Indeed, it makes no difference if the $\mathcal{DM}$ acquires information often and gradually or rarely and abruptly if he is required to learn by generating noisy data himself.

This result, nonetheless, hinges on (at least) two assumptions. First, assuming that the flow costs depend on the amount of information acquired is not without loss. For example, if the flow costs depend on a general function of the model parameters, it is not that similarly stark results can be found. Secondly, my results depend on the assumption that an experiment is a data generating process that always generates measurement error.

References

[1] Bolton, Patrick, and Christopher Harris. 1999. "Strategic signal". Econometrica (ECTA) 67 (2), 349–374.

[2] Caplin, Andrew, Mark Dean, and John Leahy. 2022. "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy." Journal of Political Economy 130 (6): 1676–1715.

[3] Bloedel, Alexander W., and Weijie Zhong. 2021. "The Cost of Optimally Acquired Information." Unpublished.

[4] Denti, Tommaso, Fabrizio Marinacci, and Aldo Rustichini. 2022. "Experimentation Cost of Information". American Economic Review (2022) 12(9): 3106–3123.

[5] Denti, Tommaso. Forthcoming. "Posterior Separable Cost of Information." American Economic Review.

[6] Denti, Tommaso, Massimo Marinacci, and Aldo Rustichini. 2022. "The experimental Order on Random Posteriors." Unpublished.

[7] Hébert, Benjamin, and Michael Woodford. 2021. "Rational Inattention When Decisions Take Time." Unpublished.

[8] Mensch, Jeffrey. 2018. "Cardinal Representations of Information." Unpublished.

[9] Morris, Stephen, and Philipp Strack. 2019. "The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs." SSRN 2991567.

[10] Øksendal , Bernt and Agnès Sulem. 2019. "Applied Stochastic Control of Jump Diffusions". Springer-Verlag Switzerland.

[11] Platten, Eckhart (1982). "An approximation method for a class of Ito processes with jump component". Lient. Mat Rink 22(2): 124-136.

[12] Pomatto, Luciano, Philipp Strack, and Omer Tamuz. 2020. "The Cost of Information." Unpublished.

[13] Mike Press. 2016. "Cleaning Big Data: Most Time-Consuming, Least Enjoyable Data Science Task, Survey Says". Forbes retrieved from: $https://www.forbes.com/sites/gilpress/2016/03/23/data-preparation-most-time-consuming-least-enjoyable-data-science-task-survey-says/amp/$

[14] Sims, Christopher A. 2003. “Implications of Rational Inattention.” Journal of Monetary Economics 50(3): 665–90.

[15] Strulovici, Bruno and Martin Szydlowski. 2015. ”On the smoothness of value functions and the existence of optimal strategies in diffusion models”. Journal of Economic Theory 159:1016–1055.

[16] Wald, A. 1947. ”Sequential analysis”. John Wiley.

[17] Zhong, Weijie. 2022 ”Optimal Dynamic Information Acquisition”. Econometrica 90(4):1537-1582.

A Proofs

A.1 Deriving dynamics.

In this section, I present the characterization of beliefs. Informally, I approximate the signal process in discrete time and take limits.

Approximating signals in discrete-time Fix some small time interval dt> 0, then an admissible signal $s = (s_t)$ can be approximated at times $t = 0, \mathrm{dt}, \ldots$ as $s_0 = 0$ and $\mathrm{d}s_t \equiv s_{t+\mathrm{dt}} - s_t$

$$\mathrm{d}s_t = \mathrm{d}_t^{\mathrm{dt}} + J_t^{\mathrm{dt}} \tag{6}$$

for $(\mathrm{d}_t^{\mathrm{dt}})_{t=0}^{\infty}$ is a sequence of independent random variables such that at time t, $d_t^{\mathrm{dt}} = \pm\sqrt{\mathrm{dt}/h_t}$ with probability $[1 - \mu_i\sqrt{h_t\mathrm{dt}}]/2$ iff $x = x_i$; meanwhile, $(J_t^{\mathrm{dt}})_{t=0}^{\infty}$ is a sequence of independent random variables such that $J_t^{\mathrm{dt}} = 1$ with probability $\lambda_{it}\mathrm{dt}$ and $J_t^{\mathrm{dt}} = 0$ with probability $1 - \lambda_{it}\mathrm{dt}$ iff $x = x_i$. This means that the independent diffusion and jump processes are weakly approximated to since the objective is to approximate a generator.

Approximating beliefs after a jump I first consider the case when there are jumps. Suppose that the $\mathcal{DM}$ held beliefs $p_{t-\mathrm{dt}} = (p_{it-\mathrm{dt}})$, then the Bayes posterior belief that

$x = x_i$ given the jump is approximately equal to

$$p_{it} = \frac{p_{it-\mathrm{dt}}\mathrm{dt}\lambda_{it}}{\sum_j p_{jt-\mathrm{dt}}\mathrm{dt}\lambda_{jt}} + o(\mathrm{dt}) = \frac{p_{it-\mathrm{dt}}\lambda_{it}}{\sum_j \lambda_{jt}p_{jt-\mathrm{dt}}} + o(\mathrm{dt})$$

where the error term $o(\mathrm{dt})$ (such that $\lim_{\mathrm{dt}\searrow 0} o(\mathrm{dt})/\mathrm{dt} = 0$) follows from the observation that distribution of d_t^{dt} approximately gives equal weight to both outcomes as dt goes to 0. Further observe that as dt goes to 0, it holds that

$$p_{it} = \frac{p_{it^-}\lambda_{it}}{\sum_j p_{jt^-}\lambda_{jt}} = \frac{p_{it^-}\lambda_{it}}{\lambda_t} = \nu_{it}.$$

It is further the case that the expected change in beliefs equals to $\mathrm{d}p_{it} \equiv p_{it} - p_{it-\mathrm{dt}}$ and satisfies that

$$\mathrm{d}p_{it} = \frac{p_{it-\mathrm{dt}}(\lambda_{it} - \lambda_t)}{\lambda_t}$$

Approximating beliefs when there is no jump Next, I characterize how beliefs change when $J_t^{\mathrm{dt}} = 0$. Suppose that the $\mathcal{DM}$ observes $\mathrm{d}s_t = \pm\sqrt{\mathrm{dt}/h_t}$, then the probability of observing said signal realization conditional on $x = x_i$, for $i = 1, 2 \ldots, n$, is

$$\mathrm{Pr}_t(\mathrm{d}s_t = \pm\sqrt{\mathrm{dt}/h_t}|x = x_i) = [1 - \lambda_{it}\mathrm{dt} \pm \mu_i\sqrt{h_t\mathrm{dt}}]/2 + o(\mathrm{dt})$$

Once again, if the prior belief is $p_{t-\mathrm{dt}} = (p_{it-\mathrm{dt}})$, then the Bayes posterior beliefs are

$$p_{it} = \frac{p_{it-\mathrm{dt}}[1 - \lambda_{it}\mathrm{dt} \pm \mu_i\sqrt{h_t\mathrm{dt}}]}{\sum_j p_{jt-\mathrm{dt}}[1 - \lambda_{jt}\mathrm{dt} \pm \mu_j\sqrt{h_t\mathrm{dt}}]} + o(\mathrm{dt}) = \frac{p_{it-\mathrm{dt}}[1 - \lambda_{it}\mathrm{dt} \pm \mu_i\sqrt{h_t\mathrm{dt}}]}{1 - \lambda_t\mathrm{dt} \pm \mu_t\sqrt{h_t\mathrm{dt}}} + o(\mathrm{dt}).$$

This equation implies that the change in beliefs $dp_{it} \equiv p_{it} - p_{it-\mathrm{dt}}$ can be approximated as

$$dp_{it} = \frac{p_{it-\mathrm{dt}}[\pm(\mu_i - \mu_t)\sqrt{h_t\mathrm{dt}} - (\lambda_{it} - \lambda_t)\mathrm{dt}]}{1 - \lambda_t\mathrm{dt} \pm \mu_t\sqrt{h_t\mathrm{dt}}} + o(\mathrm{dt})$$

and the probability that beliefs change by the amount described above occurs with a probability of approximately $(1 - \lambda_t\mathrm{dt} \pm \mu_t\sqrt{h_t\mathrm{dt}})/2$.

Approximating the expected change in beliefs Given the approximations for the change in beliefs given above, I now take expectations. First, conditional on there being no jump, the expected change in beliefs equals to

$$E[dp_{it}|J_t^{\mathrm{dt}}=0] = p_{it-\mathrm{dt}}[(\mu_i-\mu_t)\sqrt{h_t\mathrm{dt}}-(\lambda_{it}-\lambda_t)\mathrm{dt}]/2 - p_{it-\mathrm{dt}}[(\mu_i-\mu_t)\sqrt{h_t\mathrm{dt}}+(\lambda_{it}-\lambda_t)\mathrm{dt}]/2 + o(\mathrm{dt})$$
$$= -p_{it-\mathrm{dt}}(\lambda_{it}-\lambda_t)\mathrm{dt} + o(\mathrm{dt}) = -\lambda_t(\nu_{it}-p_{it-\mathrm{dt}})\mathrm{dt} + o(\mathrm{dt}).$$

Note that the expected change in beliefs when there is no jump is

$$E[dp_t|J_t^{\mathrm{dt}}=1] = \nu_t - p_t.$$

As a consequence, the unconditional change in beliefs just equals to the expectation over the conditional expectations i.e.,

$$E[dp_{it}] = \lambda_t\mathrm{dt}\left(\frac{p_{it-\mathrm{dt}}(\lambda_{it}-\lambda_t)}{\lambda_t}\right) + (1-\lambda_t\mathrm{dt})E[dp_{it}|J_t^{\mathrm{dt}}=0] = o(\mathrm{dt}).$$

Observe that this term holds due to the Law of Iterated expectations. I will exploit this observation again when estimating the change of a variable. Further notice that the approximations imply that beliefs will form a martingale.

Approximating the co-movement of beliefs Next, conditional on no jumps to co-movement of the belief that $x = x_i$ or x_j (for $i, j = 1, 2, \ldots, n$) is

$$\mathrm{d}p_{it}\mathrm{d}p_{jt} = \left[\frac{h_t p_{it-\mathrm{dt}}p_{jt-\mathrm{dt}}(\mu_i-\mu_t)(\mu_j-\mu_t)}{1\pm\mu_t^2 h_t\mathrm{dt}}\right]\mathrm{dt} + o(\mathrm{dt})$$

occurs with a probability of roughly $(1 \pm \mu_t^2 h_t\mathrm{dt})/2$. The expected covariance in the change of said observations is then equal to

$$E[\mathrm{d}p_{it}\mathrm{d}p_{jt}|J_t^{\mathrm{dt}}=0] = h_t p_{it-\mathrm{dt}}p_{jt-\mathrm{dt}}(\mu_i-\mu_t)(\mu_j-\mu_t)\mathrm{dt} + o(\mathrm{dt}).$$

Approximating the generator Lastly, I approximate the generator. Let $f : \Delta^{n-1} \to \mathrm{R}$ be a twice continuously differentiable function, then I need to calculate $E[df(p_t)]/\mathrm{dt}$ as dt goes to 0. Note that $\mathrm{d}f(p_t) = f(p_t) - f(p_{t-\mathrm{dt}})$. I can partition the expectation by the law of iterated expectations. If there is a jump, then

$$E[df(p_t)|J_t^{\mathrm{dt}}=1] = f(\nu_t) - f(p_{t-\mathrm{dt}}).$$

Alternatively, there may have been no jumps, then the change in beliefs can be approximated via a quadratic Taylor approximation as

$$\begin{aligned}E[df(p_t)|J_t^{\mathrm{dt}}=0] &= E[\nabla f(p_t)\cdot(dp_{it})+(dp_{it})Hf(p_t)\cdot(dp_{it})/2|J_t^{\mathrm{dt}}=0]+o(\mathrm{dt})\\ &= -\nabla f(p_t)\cdot(\nu_t-p_{t-\mathrm{dt}})\lambda_t\mathrm{dt}+\frac{h_t}{2}\mathrm{dt}\sum_{ij}p_{it-\mathrm{dt}}p_{jt-\mathrm{dt}}(\mu_i-\mu_t)(\mu_j-\mu_t)f_{ij}(p_{t-\mathrm{dt}})+o(\mathrm{dt})\end{aligned}$$

Since the probability of a jump is approximately $\lambda_t$dt, then the unconditional expectation equals to

$$\begin{aligned}E_t[df(p_t)] &= \lambda_t\mathrm{dt}[f(\nu_t)-f(p_t)]+[1-=\lambda_t\Delta]\times E[df(p_t)|J_t^{\mathrm{dt}}=0]\\ &= \frac{h_t}{2}\mathrm{dt}\sum_{ij}p_{it-\mathrm{dt}}p_{jt-\mathrm{dt}}(\mu_i-\mu_t)(\mu_j-\mu_t)f_{ij}(p_{t-\mathrm{dt}})+[f(\nu_t)-f(p_{t-\mathrm{dt}})-\nabla f(p_t)\cdot(\nu_t-p_{t-\mathrm{dt}})]\lambda_t\mathrm{dt}+o(\mathrm{dt})\end{aligned}$$

Dividing both sides of the expression above by dt and taking the limit as dt goes to 0 it yields that

$$\mathcal{L}f(p_t)=\frac{h_t}{2}\sum_{ij}p_{it}p_{jt}(\mu_i-\mu_t)(\mu_j-\mu_t)f_{ij}(p_t)+\lambda_t[f(\nu_t)-f(p_t)-\nabla f(p_t)\cdot(\nu_t-p_t)].$$

This concludes the proof.