Introspection

Joshua Limpert

October 2025

Abstract

A key finding in the literature on political behaviour is that attitudes are unstable across time and domain. Existing models of attitude expression fail to explain the effect of key variables like cost of thinking and cognitive uncertainty on attitude instability. We propose a model of attitude expression based on the Drift Diffusion Model. We find a closed form solution for this model and show that the model predicts temporal and domain instability. We also show that instability increases as cost of thinking goes down and that instability increases with cognitive uncertainty. In an extension, we show that our core model can be used to explain the phenomenon of priming and framing.

1 Introduction

Attitudes are unstable across time and domain. An individual asked to express their opinion on prison reform may have dramatically different answers on Tuesday compared to Thursday. And the same individual may appear liberal when asked about abortion and conservative when asked about gay rights. This is a central finding in the literature on Political Behaviour (Converse 1964; Zaller 1992).

There is significant heterogeneity in measured instability between and within people. Existing models such as the Receive - Accept - Sample model and the Gerber & Green (1993) model can explain some sources of heterogeneity but fail to explain other major sources. For example, existing models cannot explain how instability varies with cost of thinking or complexity of questions.

In this paper we propose a model of attitude expression based on the Drift Diffusion Model from psychology and Neuroscience (Swensson 1972; Luce 1986; Ratcliff & McKoon 2008; Fudenberg, Strack & Strzalecki 2018). An agent is cognitively constrained and does not know their true preferences. They can think, thinking gives them noisy (Brownian) signals of their true preferences. Thinking is costly. At a certain point, the agent stops thinking and reports an estimate of their preferences.

We find a simple, closed form solution for this model and show that the model predicts temporal and domain instability. We also show that this model can explain key heterogeneity in instability unaddressed by previous models. We show that instability increases as cost of thinking decreases. This result is surprising, but has an intuitive explanation. As people think more, there is a wider range of values that their cognitive signal can take, this increases the variance/instability of their reported preferences. We also find that instability increases as with complexity /cognitive constraint. When a topic is more complicated, the cognitive signal that people receive has higher variance, this increases the range of values that the cognitive signal can take, which increases the variance/ instability of their reported preferences.

In an extension, we show that our core model can be used to explain the phenomenon of priming and framing.

The remainder of our paper unfolds as follows. In section 2 we review the related literature. In section 3, we propose the model. In section 4, we solve the model. In sections 4 & 5 we demonstrate how this model can explain temporal and domain instability and give predictions on how cost of thinking and cognitive uncertianty affect instability. In section 6 we show how the model can explain priming and framing effects. In section 8 we show the main findings are robust to alternative specifications of the utility function and priors.

2 Contribution/Literature Review

There are a number of models of attitude formation and expression in Political Science. Achen (1992) develops a Bayesian model of socialization and attitude formation. Gerber & Green (1999) extend the standard Bayesian model to allow for certain types of bias. Neither of the papers account for response instability. Zaller (1992) proposed the Receive - Accept - Sample model. In the model, people first must understand and accept pieces of information. When they express their preferences, they randomly sample the available, accepted pieces of information to form an estimate of their preferences. The Receive-Accept-Sample model does not account for key sources of heterogeneity in temporal and domain instability of preferences. For example, it does not provide meaningful expectations about how instability changes when questions are more complicated, when people are cognitively depleted, or when people do not have enough time to properly consider questions.

The model presented in this paper is based on the idea of cognitive uncertainty. Cognitive uncertainty - subjective uncertainty about personal preferences and beliefs - has been widely documented by empirical work in psychology and behavioral economics (eg. Ilut & Valchev 2023; Gabaix 2019; Woodford 2020; Enke & Graber 2023). Enke & Graeber (2023) propose a simple framework to understand cognitive uncertainty. An agent does not know their true preferences. They see a noisy signal of their preferences and use that to form an estimate of their preferences. We enrich the Enke & Graver(2023) framework by allowing costly acquisition of signals. By thinking more an agent can acquire a more precise understanding of their true beliefs. There are several models of costly acquisition of information including models of experimentation (Keller, Rady& Cripps 2005; Cripps, Rady & Keller 2002; Callander 2011), models of rational inattention (Sims 2003) and models of sticky information (Gabaix & Laibson 2002). The model most similar to ours is the Uncertian Drift Diffusion Model proposed by Fudenberg, Strack & Strzalecki (2018). In their model an agent pays a constant per unit cost of time to observe Brownian signals of their true parameter. Two notable theoretical contributions of our paper. First, our model allows for a continuum of possible parameter values Fudenberg, Strack & Strzalecki (2018) only allow for two parameter values. Second, We find a closed form solution to the model. Fudenberg, Strack & Strzalecki (2018) only make qualitative comments about the solution.

The core theoretical contribution of this paper is a novel costly information acquisition technology with a solution that is easy to understand and easy to apply in more complex settings. This technology works particularly well in an introspection setting, but could also be applied to a vast array of other settings (Eg. policy makers searching for an ideal policy, or voters deciding between candidates).

3 The Model

An agent must report their true parameter $\theta \in R$. If their report d is inaccurate, they suffer from a dis-utility $W = -(d-\theta)^2$.

Due to cognitive constraints, the agent does not know their true parameter. They have a prior $\theta \sim N(\mu, \sigma_0^2)$ over the true parameter. They can think in order to learn more about the parameter. If they think, they receive a Brownian signal $s_t = \theta dt + \sigma dB(t)$ of their true parameter. We define ζ_t as the total filtration at time t and $S_T = \int_0^T s_t dt$. The agent must pay a cost cdt to think.

The agent chooses a level of thought τ and a report d in order to maximize their overall utility.

$$(d^*, \tau^*) = argmax_{d,\tau} E[-c\tau - (d-\theta)^2]$$

3.1 Interpretation of the Model

The following is a way to interpret key forces at play in the model. The person has a cognitive constraint which prevents them from knowing their true parameter. They can think, which gives them a noisy (Brownian) signal of their true parameter. The signal is noisy because thinking does not perfectly reveal preferences. For example, a person is asked to report their preferred tax rate. They think: "Hmm, what is my preferred tax rate?" They recall different pieces of information- conversations they had, articles they read, new segments they watched. Recollection is imperfect and somewhat random. Some of the information supports higher tax rates and some of the information supports lower tax rates; none of the information perfectly reveals their "true" preferred tax rate. They make an estimate of their "true" preferred tax rate based on the recalled information. This is a common interpretation of the process of introspection. It is the interpretation frequently used by models of dynamic information acquisition (eg. Experimentation - Cripps 2005) and by models of attitude expression (eg. Receive Accept Sample - Zaller 1992).

σ represents the noisiness of the cognitive signal or the complexity. This can be exogenously manipulated. For example, if a question is very complicated, a person will have a weaker cognitive signal of their preferences. This is an interpretation frequently used by models of cognitive uncertainty (eg. Enke & Graeber 2023).

c represents the cost of thinking. The cost of thinking may vary between people. For example, a well educated person may have a lower cost of thinking than a person with little education. The cost of thinking can also vary within a person. For example, when a person is cognitively depleted they may have a higher cost of thinking.

4 Solution to the Model

In this section we solve the model. To solve the model we first assume a fixed stopping time and find the optimal report d for the fixed stopping time. Once we obtain the optimal d we can plug that into our utility expression and determine the optimal stopping time.

Lemma 1: Optimal Report

For fixed τ

$$d^* = E[\theta|X_t] = \frac{\frac{\mu}{\sigma_0} + \frac{X_\tau}{\sigma^2}}{\sqrt{\frac{1}{\sigma^2} + \tau}}$$

This result is entirely expected. The optimal report is the expectation of θ given the information seen up to that point. The particular expression for expectation comes from the fact that S_τ is Brownian with drift. Brownian motion with drift has the nice property that $S_\tau \sim N(\theta\tau, \tau)$. This means the posterior of θ is $N(\frac{S_\tau \tau \sigma_0^2 + \mu\tau}{\sigma_0^2 + \tau}, \frac{\sigma_0^2 \tau}{\tau})$. The expectation follows.

Something to note: the optimal report is the expectation given the information at time t. In our setting, this is equivalent to the estimate of θ determined using the Kalman-Bucy filter.

Once we have the optimal report for a fixed stopping time, we can determine the expected dis-utility from the optimal report.

Lemma 2: Dis-utility

For fixed τ

$$E[W(d^*, \tau)] = -\frac{\sigma_0^2 \sigma^2}{\sigma^2 + \tau\sigma_0^2}$$

$E[W(d^*, \tau)] = E[(d^* - \theta)^2]$ is the mean squared error of the optimal report d^*. Mean squared error can be decomposed into bias and variance. In Lemma 1 we found that d^* is the unbiased estimate of θ. This means $E[W(d^*, \tau)]$ is the variance of the estimate d^*. As we noted previously, d^* is the estimate of θ determined by the Kalman - Bucy filter.$E[W(d^*, \tau)]$ is the variance of the Kalman-Bucy estimate of θ.

$E[W(d^*, \tau)]$ is the expected dis utility from the optimal report. The expected dis-utility decreases when a person thinks more (τ increases), and the expected dis-utility increases when the complexity or cognitive uncertainty increases (σ^2 increases).

Thinking more reduces expected dis-utility. The optimizing agent wants to balance the cost of thinking more with the benefit. The optimal stopping rule captures this fact:

Proposition 1: Optimal Action

$$\tau^* = \begin{cases} \frac{\sigma}{\sqrt{c}} - \frac{\sigma^2}{\sigma_0^2} & \text{if } \sqrt{c} < \frac{\sigma_0^2}{\sigma} \\ 0 & \text{otherwise} \end{cases}$$

$$d^* = \begin{cases} \frac{\sqrt{c}}{\sigma} S_{\tau^*} + \mu \frac{\sqrt{c}\sigma}{\sigma_0^2} & \text{if } \sqrt{c} < \frac{\sigma_0^2}{\sigma} \\ 0 & \text{otherwise} \end{cases}$$

To find τ^* we took first order conditions. The optimal stopping time is the precise time when the marginal cost of thinking exceeds the marginal benefit. To find d^*,we plugged τ^* into the optimal reporting rule determined in lemma 1. This is a closed form solution to an Uncertain Drift Diffusion Model. Previous work on uncertain drift diffusion models have not obtained closed form solutions. They are only able to qualitatively describe elements of the solution (eg. Fudenberg, Strack & Strzalecki 2018).

Proposition 2: Comparative Statics of Optimal Stopping

- $\frac{\partial \tau^*}{\partial c} < 0$
- $\frac{\partial \tau^*}{\partial \sigma_0^2} > 0$
- $\frac{\partial \tau^*}{\partial S_\tau} = 0$

The amount of time spent thinking increases when the cost of thinking decreases (Item 1). This is entirely expected. Thinking more increases utility, when it is cheaper to think, people will think more.

The amount of time spent thinking increases when the amount of complexity or cognitive uncertainty increases (Item 2). This is also entirely expected. When there is high complexity or cognitive uncertainty a person needs to think more to achieve the same level of understanding about their preferences.

There is an unexpected result- the amount of time spent thinking does not depend on the history of observations (Item 3). To our knowledge, this is the only model of dynamic costly information acquisition with a flexible stopping rule where the amount of information collected is not affected by the history. For example, models of experimentation have history dependent stopping rules (eg. Kelly Radner & Crips 2005). Why do we get this result? $E[W(d^*, \tau)] = E[(d - \tau)^2]$ is effectively the mean squared error of the report d. Mean squared error can be decomposed into bias and variance. The report is unbiased. The amount by which additional data shrinks the variance does not depend on the history of observations. In extensions we examine whether this finding is robust to different forms of $W(d^*, \tau)$.

Proposition 3: Attenuation

- $\frac{\partial E[d^*]}{\partial \theta} < 1$
- $\frac{\partial^2 E[d^*]}{\partial \theta \partial \sigma^2} < 0$
- $\frac{\partial^2 E[d^*]}{\partial \theta \partial c} < 0$

This model finds attenuation (item 1). Attenuation - the fact that reported references are insufficiently elastic to true preferences - is one of the central findings in the cognitive uncertainty literature (eg. Enke & Graeber 2021). Attenuation occurs because people form noisy estimates of the true parameter θ. To form the estimates they combine information from their noisy signal with information from their prior. This estimation technique pulls reports towards the mean of the prior which makes reports insufficiently elastic to the true parameter.

We find that attenuation increases when there higher cognitive uncertainty(item 2) . When there is higher cognitive uncertainty, uncertainty people weight their prior information more heavily, this means they are less responsive to information from thinking. This is consistent with empirical work and existing models of cognitive uncertainty.

We find that attenuation increases when there are higher cognitive costs. When costs are higher, a person thinks less. When they think less they have less information about their true preferences and they rely on their prior more. Existing models of cognitive uncertainty do not capture this fact.

5 Application 1 - Temporal Instability

In the next two section we discuss two applications- Temporal instability and Domain instability. The model predicts both temporal and domain instability. The model also gives precise predictions about how cost of thinking and complexity affect temporal and domain instability.

We model temporal instability in the following way. An agent needs to report the same parameter θ at two distinct times, in period 1 and in period 2. In each period they conduct the introspection process described in the main model. We make two assumptions. First, in period 2, they cannot remember anything from period 1. Second, the Brownian processes in period 1 and period 2 are independent. We define the variable $D = d_1 - d_2$.

Proposition 4: Temporal Instability

$$D \sim N(0, 2\sigma^2(1 - \frac{\sqrt{c}}{\sigma_0^2})^2)$$

Two things to note about this result. First, the difference in the reports has a mean of zero. Period 2 reports are not expected to be systematically higher or lower than period 1 reports. Second, the variance is greater than zero. Period 1 and Period 2 reports are not necessarily the same. This occurs because people form estimates of their true preferences from noisy introspective signals. The signals, may be different in period 1 and period 2. This gives us temporal instability. Temporal instability is one of the central findings of political behaviour, any model attitude expression should yield temporal instability.

This model gives predictions about how key factors like cost of thinking and cognitive uncertainty affect temporal instability.

Corollary 1: Comparative Statics of Temporal Instability

- $\frac{\partial Var(d^*)}{\partial c} > 0$
- $\frac{\partial Var(d^*)}{\partial \sigma^2} < 0$

We expect temporal instability to fall when the cost of thinking rises (item 1). This is a very surprising result, but there is an intuitive explanation. When cost of thinking increases, the amount of time people spend thinking increases, the range of possible values of the cognitive signal S_{τ^*} is larger. This means the difference between the report at time one and time two can be larger. In other words, temporal instability is higher. This result highlights the benefit of studying attitude expression with a formal model, it is not clear whether we would have obtained this result using intuition alone. When does cost of thinking rise? An example is fatigue. From this model, we expect temporal instability to fall when a person is more fatigued. Existing models of attitude expression do not make predictions about response instability and fatigue. To our knowledge there is no empirical work in political science that examines temporal instability and fatigue.

We expect temporal instability to rise when a person has more cognitive uncertainty about the parameter θ (item 2). When the cognitive uncertainty about the parameter is higher, the cognitive signal has more variance, the range of possible values of the cognitive signal is larger and the range of possible expressed attitudes is larger. An example of a treatment that would increase cognitive uncertainty is complicated questions. Existing models of attitude expression to not make predictions about response instability and question complexity. To our knowledge there has been no work in Political Science on response instability and question complexity.

6 Application 2 - Domain Instability

We model domain instability in the following way. An agent is asked to report their preferences between two domains θ_a, θ_b. Across the population the true covariance between θ_a and θ_b is $cov(\theta_a, \theta_b) = \rho > 0$. For each parameter, they conduct the introspection process described by the model. We make two assumptions. First, for domain a, they cannot remember their report or introspection process for domain b and vise versa. Second, the Brownian processes for a and b are independent. We define $\beta = \frac{cov(\theta_a,\theta_b)}{var(\theta_a)var(\theta_b)}$ and we define $\hat{\beta} = \frac{cov(d_a,d_b)}{var(d_a)var(d_b)}$.

Proposition 5: Domain Instability

$$\hat{\beta} = \frac{var(\theta_a)var(\theta_b)}{(var(\theta_a) + 2\sigma^2(1-\frac{\sqrt{c}}{\sigma_0^2})^2)(var(\theta_b) + 2\sigma^2(1-\frac{\sqrt{c}}{\sigma_0^2})^2)}\beta < \beta$$

Cognitive noise increases the variance of reported parameters, this lowers the estimated relationship between parameters. Cognitive uncertainty makes it appear as if there is low cross domain constraint (stability). Our model makes specific predictions about how variables like cost of thinking and cognitive uncertainty affect apparent cross domain instability.

Corollary 1: Comparative Statics of Domain Instability

- $\frac{\partial \hat{\beta}}{\partial c} < 0$
- $\frac{\partial \hat{\beta}}{\partial \sigma_0^2} < 0$

When cost of thinking decreases, the measured relationship between θ_a and θ_b will decreases (item 1). When people have a lower cost of thinking, they think more, this increases the variance of d^* and the measured relationship between θ_a and θ_b will be lower. It will appear as if there is high cross domain instability.

When there is higher cognitive uncertainty about the parameters, the measured relationship will decrease (item 2) When there is higher cognitive uncertainty about the parameters θ_a and θ_b, there will be higher variance of d^*, this lowers the measure relationship between θ_a and θ_b.

There have been useful tools developed by Political Methodologist and Econometricians to address domain instability. For example, the ORIV estimator (Perez et al 2021) is designed to correct artificially low measured relationships between variables due to measurement or response errors (like cognitive uncertianty). It would be an interesting project to measure true domain constraint using an ORIV estimator.

7 Application 3 - Priming and Framing

Priming and Framing - the fact that calling attention to certain facts or aspects of an issue changes expressed opinions - has been studied widely in political science (eg. Druckman 2007;

Druckman et al. 2001). We can use this model to explain the phenomena of priming and framing.

We model it in the following way. An attitude $\hat{\theta}$ is comprised of two parts. Domain 1: θ_1 and domain 2: θ_2. In particular $\hat{\theta} = \alpha\theta_1 + (1-\alpha)\theta_2$. The person is asked to report $\hat{\theta}$. We assume they know α, but do not know θ_1 or θ_2. Initially they have the same cost of thinking for domain 1 and domain 2, $c_1 = c_2$. Priming or framing reduces the cost of thinking in a particular domain.

Priming occurs in domain 1, making $c_1 < c_2$. For ease of exposition we assume that $\theta_1 > \mu > \theta_2$. We also assume that the Brownian processes of domain 1 and domain 2 are independent.

$\hat{d}_p$ is the report of $\hat{\theta}$ with priming. $\hat{d}_{np}$ is the report of $\hat{\theta}$ with no priming.

Proposition 6: Priming

$$E[\hat{d}_p] < E[\hat{d}_{np}]$$

To develop an estimate of $\hat{\theta}$, the person must develop two separate estimates one of θ_1 and one of θ_2. Priming lowers the cost of thinking about domain 1, this means the agent thinks more about the domain. As they think more about the domain, they put less weight on the prior and more weight on the realized signals. Because $\theta_1 > \mu$, the expectation of the estimate is higher. The person has been primed to think about domain 1, the expected estimate will be closer to θ_1.

We should note this is very different from other models of priming and framing. Other models of priming and framing (eg. Schliefer 1993; Ortelova 2025) conceptualize priming and framing as changing the weight placed on each domain. This would be equivalent to changing α. The old models are valuable, but it is clear that cognitive complexity may also play a role in priming and framing.

8 Robustness Checks

In this section we examine alternative specifications of the utility function, we show that most of our main results continue to hold.

8.1 Dirac Utility

Suppose the utility function is now $W = \begin{cases} \frac{1}{2\epsilon} & \text{if } d \in (\theta+\epsilon, \theta-\epsilon) \\ 0 & \text{otherwise} \end{cases}$. With this utility function agents get a fixed payoff for being close enough to the true parameter and a payoff of zero for not being close enough to the true parameter.

We continue to assume that the agent has a prior $\theta \sim N(\mu, \sigma_0^2)$ over the true parameter and that they can think to learn more about the parameter. If they think, they receive a Brownian signal $s_t = \theta dt + dB(t)$ of their true parameter. The agent must pay a cost cdt to think.

To solve this model we again find the optimal report d^* for a fixed stopping time τ. We then plug the reporting rule into the expression for utility and take first order conditions to find the optimal stopping time.

Lemma 3: Optimal Report - Dirac Dis-utility

For fixed τ

$$d^* = E[\theta|\zeta_t] = \frac{S_\tau \tau \sigma_0^2 + \mu\tau}{\sigma_0^2 + \tau}$$

Again in this case, the optimal report is the posterior mean/mode. The agent is most likely to land within ϵ of the true parameter and receive positive utility if they report the mode of the posterior distribution. In the case of normal distributions the mean and the mode co-coincide.

Lemma 4: Dirac Utility

For fixed τ

$$E[W(d^*, \tau)] = \sqrt{\frac{\sigma_0^{-2} + \tau}{2\pi}}$$

This result was obtained from a bunch of algebra, details are in the appendix. This expression makes sense, utility is increase in the amount of time spent thinking and decreasing in the amount of initial uncertainty about the parameter.

Thinking more gives the agent more utility, there is a cost to thinking more. The optimizing agent balances the cost with the benefit.

Proposition 7: Optimal Action

$$\tau^* = \begin{cases} \frac{1}{8\pi\sigma_0^2} - \frac{1}{\sigma_0^2} & \text{if } 8\pi\sigma_0^2 < \sigma_0^2 \\ 0 & \text{otherwise} \end{cases}$$

$$d^* = \begin{cases} 8\pi c^2 S_{\tau^*} + \mu \frac{8\pi c^2}{\sigma_0^2} & \text{if } \sqrt{c} < \sigma_0^2 \\ \mu & \text{otherwise} \end{cases}$$

Again the optimal stopping time τ^* was determined by taking first order conditions, and the optimal reporting rule was determined by plugging the optimal stopping time into the expression for the optimal reporting rule.

There are several similarities between this specification of the model and the main model. First, the stopping rule does not depend on the history. Second, the optimal stopping time is decreasing in c and increasing in σ_0^2. Third, the model predicts attenuation.

8.2 Chernoff Utility

Suppose the utility from a report was given by $W = \begin{cases} k & \text{if } sign(d) = sign(\theta) \\ -k & \text{otherwise} \end{cases}$. With this utility function agents get a positive payoff for guessing the correct sign of their parameter. For example, they get a positive payoff for guessing that they prefer higher taxes opposed to lower taxes. The precise tax rate doesn't matter.

Again we assume that the agent has a prior $\theta \sim N(\mu, \sigma_0^2)$. That thinking , results in Brownian signals $s_t = \theta dt + dB(t)$ of their true parameter. And the agent must pay a cost cdt to think.

Once again we start by fixing τ and finding the optimal reporting rule.

Lemma 5: Optimal Report - Dirac Dis-utility

For fixed τ

$$d^* = \begin{cases} + & \text{if } S_\tau + \frac{\mu}{\sigma_0^2} > 0 \\ - & \text{otherwise} \end{cases}$$

Here the agent reports that the parameter is positive if the probability that the parameter is positive is greater than $\frac{1}{2}$. For a fixed stopping time τ, this occurs when $S_\tau + \frac{\mu}{\sigma_0^2} > 0$.

Using this decision rule we can compute the expected value of W at time τ.

Lemma 6: Chernoff Utility

$$E[W(d^*, \tau)] = \frac{1}{\sqrt{\tau + \sigma_0^{-2}}} \phi\left(\frac{S_\tau + \frac{\mu}{\sigma_0^2}}{\sqrt{\tau + \sigma_0^{-2}}}\right) - \frac{|S_\tau + \frac{\mu}{\sigma_0^2}|}{\sqrt{\tau + \sigma_0^{-2}}} \Phi\left(\frac{|S_\tau + \frac{\mu}{\sigma_0^2}|}{\sqrt{\tau + \sigma_0^{-2}}}\right)$$

This result was determined by (Shiryaev et al 2008). As part of their work on Chernoff Problems. They come up with an integral equation which defines the optimal stopping time. We come up with a closed form solution for this problem. The following lemmas are necessary intermediate steps for the solution of this problem. We can redefine $\tau^* = T + \tau_r^*$. T is the time spent searching so far and τ_r^* is the time remaining until the agent reaches τ^*.

Lemma 7: Optimal Stopping Time

$$\frac{\partial}{\partial \tau_r^*} E[c\tau + W(d^*, \tau^*)] = 0$$

The derivative of the utility, with respect to τ_r^* is zero at the optimal stopping time. This is by definition. τ^* is designed to maximizes expected utility, the derivative must be zero.

Lemma 8: Optimal Stopping Time

$$\tau_r^* = 0 \ @\tau^*$$

At the optimal stopping time, there is no more time remaining until reaching the optimal stopping time. This again is by definition. We can combine Lemma 7 and 8 to derive the optimal stopping time.

Proposition 8: Optimal Action

This proposition is a simple solution to this optimal stopping problem. This optimal stopping problem was first proposed by Chernoff in 1964. Existing solutions require an integral equation to define the stopping rule (eg. Gapeev & Peskir 2006; Ekström & Vaicenavicius, 2015). Our solution does not require an integral equation, the proof technique is also much simpler. This is one of the theoretical contributions of the paper.

We know the derivative is zero at the optimal stopping time. We also know that the additional time left to reach the optimal stopping time is zero. We can plug zero into the equation, this defines t in terms of S_τ.

In parameterization of the problem, the stopping time does depend on the history. This makes sense. If the cumulative signal S_t is close to zero, the agent will want to continue thinking; however, if the cumulative signal is far from zero, there is little benefit to continuing thinking. The stopping rule depends on t. When t is low, the agent will stop thinking if S_t is relatively close to 0. When t is large, the cumulative signal needs to be further from 0 for the agent to stop thinking. This makes sense, the variance of the cumulative signal grows with t, the a positive cumulative signal early means a higher likelihood that the true parameter is positive, compared to the same positive cumulative signal later on.

9 Bibliography

Achen, C.H., 1992. Social psychology, demographic variables, and linear regression: Breaking the iron triangle in voting research. Political behavior, 14(3), pp.195-211.

Callander, S., 2011. Searching and learning by trial and error. American Economic Review, 101(6), pp.2277-2308.

Converse, P.E., 2006. The nature of belief systems in mass publics (1964). Critical review, 18(1-3), pp.1-74.

Cripps, M., Keller, G. and Rady, S., 2002. Strategic experimentation: The case of Poisson bandits. Available at SSRN 318778.

Enke, B. and Graeber, T., 2023. Cognitive uncertainty. The Quarterly Journal of Economics, 138(4), pp.2021-2067.

Fudenberg, D., Strack, P. and Strzalecki, T., 2018. Speed, accuracy, and the optimal timing of choices. American Economic Review, 108(12), pp.3651-3684.

Gabaix, X., 2019. Behavioral inattention. In Handbook of behavioral economics: Applications and foundations 1 (Vol. 2, pp. 261-343). North-Holland.

Gabaix, X., Laibson, D. (2005). Bounded rationality and directed cognition. Harvard University, 20.

Gerber, A. and Green, D., 1999. Misperceptions about perceptual bias. Annual review of political science, 2(1), pp.189-210.

Ilut, C. and Valchev, R., 2023. Economic agents as imperfect problem solvers. The Quarterly Journal of Economics, 138(1), pp.313-362.

Keller, G.; Rady, S. and Cripps, M., 2005, 'Strategic Experimentation with Exponential Bandits,' Econometrica, 73, No. 1, pp. 39–68.

Luce, R. Duncan. 1986. Response Times: Their Role in Inferring Elementary Mental Organization. New York: Oxford University Press.

Ratcliff, R. and McKoon, G., 2008. The diffusion decision model: theory and data for two-choice decision tasks. Neural computation, 20(4), pp.873-922.

Sims, C.A., 2006. Rational inattention: Beyond the linear-quadratic case. American Economic Review, 96(2), pp.158-163.

Swensson, R.G., 1972. The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. Perception Psychophysics, 12(1), pp.16-32.

Woodford, M., 2020. Modeling imprecision in perception, valuation, and choice. Annual Review of Economics, 12(1), pp.579-601.

Zaller, J., 1992. The nature and origins of mass opinion. Cambridge university press.