The Uncertainty Trap: Thresholds, Noise, and Time Preference Inconsistency *

Inhwa Kim†
Harris-Stowe State University

Abstract

This paper introduces a dynamic threshold model that provides a comprehensive theoretical framework for understanding time preferences by incorporating stochastic shocks, adaptive persistence, and recursive subjective thresholds. The model advances beyond traditional static approaches, such as the β-δ framework, emphasizing the dynamic and context-sensitive nature of decision-making under uncertainty. The model's recursive structure captures behavioral instability near subjective thresholds, long-term path dependence caused by accumulated stochastic shocks, and the trade-off between stability and responsiveness in adaptive environments. The model elucidates how heightened uncertainty and stochastic noise amplify preference reversals near decision thresholds, providing a structural explanation for observed behavioral volatility. Over the long term, cumulative stochastic shocks embed persistent instability into preference trajectories, creating path dependence that prevents reversion to stable baselines. The adaptive persistence parameter emerges as a mechanism for balancing resilience to volatility with responsiveness to external changes. By addressing both short-term fluctuations and long-term fragility, the dynamic threshold model bridges limitations in traditional frameworks. Its theoretical propositions align with empirical findings, offering insight into phenomena such as cognitive processing delays, behavioral instability near uncertainty thresholds, and enduring effects of environmental shocks. Beyond theoretical contributions, the model offers insights for experimental designs, behavioral interventions, and policy frameworks aimed at stabilizing preferences in uncertain decision contexts. The dynamic threshold model provides a flexible foundation for advancing the understanding of time preferences, offering both theoretical rigor and practical relevance.

Keywords: dynamic threshold; stochastic shocks; time preference; behavioral instability
JEL Classification Codes: C60, C61, D90, D91

*All remaining errors are my own.

†Corresponding author. Anheuser-Busch School of Business, Harris-Stowe State University, Saint Louis, MO 63103, USA. Email: kimi@hssu.edu.

1 Introduction

The process of intertemporal decision-making, wherein individuals evaluate trade-offs between outcomes realized at different points in time, represents a cornerstone of economic and behavioral inquiry (Samuelson, 1937). Traditional models of this process, notably exponential discounting, often assume a consistent rate of discounting future utilities, thereby positing temporal stability in individual preferences. This assumption, however, has been challenged by a considerable body of empirical research. Such studies have demonstrated that preference structures are frequently characterized by time-inconsistency, evidenced by phenomena such as present bias, magnitude effects, and sub-additive discounting (Frederick, Loewenstein, & O'Donoghue, 2002; Laibson, 1997; Read, 2001). These findings suggest that standard functional forms may not fully capture the intricate nature of temporal decision-making, thus highlighting the need for alternative theoretical frameworks.

Models of time preference are central to understanding intertemporal choice; however, capturing the full complexity of observed human behavior remains challenging. In response to the limitations of exponential discounting, particularly its inability to explain time-inconsistent behaviors, the β-δ model emerged as a prominent framework (Laibson, 1997). By introducing the present-bias parameter β, this model accounts for the disproportionate weighting of immediate rewards relative to future rewards, while the δ parameter captures the exponential discounting of future utilities. Similarly, hyperbolic discounting assumes that discount rates decline over time without introducing a specific parameter for present bias. Both approaches address inconsistencies that static exponential discounting cannot. However, these models do not account for the role of external stochastic forces or environmental variability in shaping time preferences.

Emerging evidence indicates that preferences are influenced not only by internal biases but also by external stochastic forces and environmental variability. Research has shown that stochastic noise affects choice consistency, making decisions unpredictable even when preferences appear stable (Kahneman, Sibony, & Sunstein, 2021; He & Bhatia, 2018). Neuroscientific findings further support the idea that both external and internal stochastic factors alter choice behaviors. Additionally, studies on attentional mechanisms demonstrate that the framing of choices and external cues significantly contribute to preference reversals (Zhou, Li, & Liu, 2021). These findings highlight the limitations of models that assume static internal parameters and underscore the need for a framework that incorporates the dynamic and stochastic nature of preferences, as well as contextual influences on decision-making (Cohen et al., 2020).

This paper builds on the rich literature on time-inconsistent preferences, particularly models that use internal benchmarks to explain intertemporal choices. It shares a foundation with the β-δ framework (Akerlof, 1991; Laibson, 1997; O'Donoghue & Rabin, 1999) and incorporates insights from recent work, such as Laibson and Maxted (2023), which demonstrates the sensitivity of present-biased models to parameter selection. Motivated by these findings, this paper moves beyond methodological considerations to explore the dynamic properties of internal benchmarks, such as subjective discount rates and uncertainty thresholds.

To address the limitations of existing models, particularly their inability to account for the dynamic effects of external stochastic forces, this paper introduces a dynamic subjective threshold (θ_t) within the β-δ discounting framework. Unlike static parameters, this threshold evolves recursively, influenced by past uncertainty, stochastic shocks, and cognitive inertia. By incorporating noise variance (σ^2), the model explains behavioral instability, particularly in contexts where uncertainty fluctuates near the threshold. This approach captures the stochastic and path-dependent nature of decision-making, providing insights into phenomena such as preference reversals and the role of environmental cues in shaping decision trajectories.

A core feature of the model is the persistence parameter, which highlights how past choices influence present decisions and demonstrates that internal parameters are not fixed but evolve based on experience and learning. This dynamic perspective explains time-inconsistent behaviors, such as present bias and loss aversion, as emergent features of interactions between decision-makers and their environment. By introducing mechanisms for dynamic changes in preferences, the model builds upon prior work, such as Harris and Laibson (2001) and O'Donoghue and Rabin (2001), and accounts for previously unexplained empirical phenomena.

Additionally, the model incorporates stochastic external forces that cause deviations from ideal values, creating instability in choice behaviors. Read (2001) highlighted that decisions often deviate from pre-set functions, illustrating the unpredictability of human choices. Building on this observation, this model provides a framework for understanding such deviations as outcomes of dynamic interactions between internal biases and external environmental factors. By explicitly integrating these dynamics, the model offers a more nuanced perspective on the mechanisms driving these inconsistencies.

The model further explores the importance of these mechanisms through a series of six testable propositions. These propositions demonstrate how choices are made at a given point in time, how they are influenced by a subjective threshold, and how such choices may change over time. The propositions also address how individuals may act to improve long term decision making. By explicitly considering these factors, the model provides a comprehensive framework for understanding how preferences change over time, and how those changes may be influenced by both the external environment, and an individuals response to that environment. This multi-faceted approach addresses limitations of other models, which often have precise mathematical structures that highlight internal factors, but ignore external influences. In contrast, this model explicitly emphasizes the interaction of both internal and external mechanisms. The paper will then go on to demonstrate how this theoretical framework compares with existing models in the literature, and will highlight the ability of this framework to account for previously observed empirical results.

This research offers new insights into the nature of choice, by proposing a theoretical model that explicitly incorporates these dynamic external forces, as well as the strategic behaviors that individuals may adopt when choosing their parameters. The results also demonstrate that preference changes over time are not a result of a fixed bias, but rather an active and dynamic process, shaped by past and present environmental conditions. By highlighting how these factors inter-

act, the aim is to provide a more nuanced framework that may be used to better understand the complexities of time preference and also highlight mechanisms that have been previously ignored by existing models. This emphasis on external factors also contrasts with the work of Halevy (2015), which provides a framework for analyzing time-consistent behaviors, but does not explicitly include stochasticity or external forces. These findings highlight the need for an approach that can address both the internal and external mechanisms that influence preferences over time. The novel framework proposed in this paper directly addresses these limitations, by providing a model that is both internally consistent, and also able to account for the dynamic and stochastic external environment.

The subsequent sections are structured as follows. The next section presents a review of the relevant literature, situating this study within the broader theoretical and empirical context. Following this, the dynamic threshold and uncertainty model is introduced, accompanied by theoretical propositions derived from its structure. Empirical evidence supporting these propositions is then discussed, followed by an analysis of the model's broader implications and limitations. The final section concludes with directions for future research.

2 Literature Review

The study of time preferences, which examines how individuals value immediate versus delayed rewards, has long been a focal point in economics and behavioral sciences. Samuelson (1937) laid the groundwork for intertemporal choice theory by introducing the exponential discounting model, where individuals discount future utilities at a constant rate. This framework assumes stable preferences over time and has served as a benchmark for subsequent theories. However, empirical anomalies, such as preference reversals and inconsistent discounting behavior, challenged this static perspective, leading to the development of alternative models.

A seminal work by Akerlof (1991) offered an early examination of time-inconsistent behavior, using the framework of procrastination and obedience to explore how individuals may struggle to act in accordance with their long-term goals. This work provides an important foundation for subsequent models that have sought to formalize and extend these ideas, including the present-biased model that is considered in this study.

In response to these inconsistencies, Laibson (1997) introduced the hyperbolic discounting model, widely referred to as the beta-delta model. This approach incorporates present bias, illustrating that individuals display higher impatience for short-term rewards while exhibiting more consistent discounting behavior for long-term rewards. The beta-delta model captures the tendency for individuals to disproportionately value immediate gratification, addressing key anomalies observed in earlier studies.

Loewenstein and Prelec (1992) provided crucial empirical evidence of time-inconsistent preferences, demonstrating that observed behaviors often deviate from the static predictions of exponential discounting. This research has played a pivotal role in highlighting the limitations of the

exponential model, and in motivating alternative frameworks such as the beta-delta and hyperbolic models.

Building on these foundational contributions, Read (2001) proposed that time discounting may not adhere strictly to an exponential or hyperbolic functional form. Instead, discounting behaviors may vary depending on the framing of decisions and specific contextual cues. This perspective emphasized the flexibility and inconsistency in observed time preferences, suggesting that individual behavior cannot be fully explained by static discounting functions alone.

Empirical studies have since reinforced these theoretical advancements by highlighting systematic deviations from traditional discounting models. Mitchell, Wilson, and Karalunas (2015) compared hyperbolic, delay-amount sensitivity, and present-bias models of delay discounting, revealing that each model captures different aspects of temporal preferences. Their findings demonstrated that no single model could comprehensively account for the dynamic nature of time preferences across diverse contexts and tasks.

The growing body of evidence on intertemporal choice has also underscored the role of cognitive mechanisms and neural processes in shaping time preferences. Lee (2013) examined the neurobiological underpinnings of decision-making, demonstrating how brain regions associated with reward processing and cognitive control interact during intertemporal choices. This neuroeconomic perspective provided further insights into the variability observed in discounting behavior, linking time preferences to underlying biological and cognitive processes.

As research progressed, it became evident that time preferences are not only influenced by intrinsic cognitive biases but also by stochastic and environmental factors. Studies by Kahneman, Sibony, and Sunstein (2021) and He and Bhatia (2018) emphasized that noise, defined as random variability in judgments, plays a critical role in shaping time preferences. These works highlighted that preference reversals are not solely a result of intrinsic psychological biases but are also driven by exogenous stochastic shocks and environmental uncertainty.

Despite these theoretical advancements, numerous empirical findings reveal substantial variability and inconsistency in time preferences across different contexts and individuals (Frederick, Loewenstein, & O'Donoghue, 2002). These inconsistencies have been attributed to both internal cognitive biases and external stochastic influences, highlighting the need for dynamic models that can accommodate these variations. Augenblick, Niederle, and Sprenger (2015) demonstrate that dynamic inconsistency persists in real effort tasks, suggesting that time preferences are not only context-dependent but also shaped by task-specific cognitive mechanisms. Similarly, Lempert and Phelps (2016) emphasize the malleability of intertemporal choices, showing that external factors, such as framing effects and emotional states, significantly alter discount rates.

Recent research underscores the role of noise and stochasticity in shaping time preferences. Kahneman, Sibony, and Sunstein (2021) argue that noise, defined as random variability in judgments, is a pervasive flaw in human decision-making processes. This perspective aligns with the findings of He and Bhatia (2018), who highlight the stochastic variability in intertemporal choices and demonstrate that observed inconsistencies often arise from random fluctuations in perceived

value and risk. These results suggest that preference reversals are not solely a result of intrinsic psychological biases but are also driven by exogenous noise and uncertainty in the decision environment. Alós-Ferrer, Jaudas, and Ritschel (2021) further illustrate that attention allocation significantly influences preference reversals, demonstrating that shifts in visual focus between attributes, such as delay and reward magnitude, can alter decision outcomes.

The relationship between attention and time preferences has been extensively explored in experimental settings. Ashby, Dickert, and Glöckner (2012) reveal that attention allocation mediates subjective valuations of rewards in consumer choices, while Ashby, Walasek, and Glöckner (2015) extend this analysis to risky decision-making environments. Attention, therefore, serves as a critical cognitive mechanism linking observed choices to underlying preferences. Stewart, Hermens, and Matthews (2015) support this view by showing that eye-tracking data can predict shifts in risk preferences, reinforcing the argument that changes in attentional focus act as a source of behavioral noise. In the domain of intertemporal choices, Franco–Watkins, Mattson, and Jackson (2016) and Sui, Liu, and Rao (2020) demonstrate that attentional processes systematically influence temporal discounting, with individuals disproportionately focusing on either delay or reward attributes depending on task framing and environmental cues.

Cognitive biases and subjective uncertainty further complicate the interpretation of time preference data. Skylark, Farmer, and Bahemia (2021) argue that intertemporal choices are heavily influenced by inference mechanisms, where individuals rely on heuristics to estimate future outcomes. This reliance introduces systematic noise into decision-making processes, particularly under conditions of incomplete information. Oberrauch and Kaiser (2022) expand on this argument by examining the role of cognitive ability and financial literacy in shaping time preferences, demonstrating that individuals with higher cognitive capacities exhibit more consistent discounting behavior. These findings suggest that cognitive noise, arising from information processing limitations, interacts with external stochastic factors to produce observed variability in time preferences.

The interaction between risk and time preferences has also received substantial attention in the literature. Somasundaram and Eli (2022) highlight that sensitivity to time delays is influenced by the probability of receiving rewards, showing that risk preferences and time preferences are not independent dimensions but interact in complex ways. This interaction challenges traditional models that treat these preferences as separate constructs. Similarly, Blavatskyy (2022) introduces a model where intertemporal choices are framed as a trade-off between cumulative payoff and average delay, providing a novel perspective on how individuals integrate probabilistic outcomes into temporal decisions.

Economic constraints and environmental factors further shape intertemporal preferences. Carvalho, Meier, and Wang (2016) provide evidence that liquidity constraints and financial shocks influence discount rates, demonstrating that economic hardship can amplify impatience and preference reversals. DellaVigna (2018) emphasizes the importance of incorporating behavioral biases and external constraints into structural economic models, arguing that real-world decision

environments often distort theoretically predicted discount rates. These findings highlight the importance of accounting for environmental variability and economic constraints when analyzing intertemporal preferences.

From a methodological perspective, noisy experimental results remain a critical challenge in estimating time preference parameters. Andersson et al. (2016) demonstrate that estimation results can be systematically biased if experimental data quality is correlated with cognitive capacity, literacy, or attentional biases. Schildberg-Hörisch (2018) reinforces this argument, suggesting that systematic noise in experimental designs may obscure underlying behavioral patterns and produce inconsistent results. These methodological challenges emphasize the need for robust experimental designs that can isolate the effects of noise, stochasticity, and cognitive biases on time preferences.

Recent research such as that of Laibson and Maxted (2023) has focused on the methodological rigor that is necessary to produce robust numerical simulation when working with present biased consumers. In particular, they show that time step specification plays a significant role in the stability of such models, and they use this to develop a new metric for determining the reliability of numerical outputs.

Overall, the existing literature demonstrates that time preferences are not static constructs but are shaped by dynamic interactions between cognitive processes, stochastic noise, attentional mechanisms, and environmental factors. Foundational theories, such as Samuelson's (1937) exponential discounting model and Laibson's (1997) hyperbolic discounting model, provide essential frameworks for understanding time preferences. However, empirical findings reveal substantial variability, highlighting the need for dynamic models capable of integrating noise, adaptive mechanisms, and recursive feedback processes into preference formation. The ongoing research underscores the complexity of intertemporal choice, pointing towards the necessity of theoretical frameworks that can capture both the structural and stochastic dimensions of time preference dynamics.

3 Theoretical Framework

This study integrates a dynamic subjective threshold mechanism into the traditional β-δ discounting model to analyze time preference inconsistency under evolving uncertainty. The β-δ model, established by Laibson (1997) and extended by O'Donoghue and Rabin (1999), provides a foundation for understanding time-inconsistent preferences. Present bias, represented by β, reflects the disproportionate weighting of immediate rewards, while the exponential discount factor, δ, captures the consistent depreciation of future utility. While effective at explaining intrapersonal conflict, these models generally overlook how subjective perceptions of uncertainty (z) influence agents' decisions to prioritize either short-term or long-term outcomes.

To address this gap, the framework introduces an evolving subjective uncertainty threshold (θ_t), which acts as a critical boundary separating two behavioral regimes. At each point in time,

agents evaluate their preferences based on their subjective perception of uncertainty relative to this threshold. When perceived uncertainty exceeds the threshold ($z > \theta_t$), agents prioritize short-term rewards (R_S) to minimize risk exposure. In contrast, when uncertainty remains below the threshold ($z \leq \theta_t$), agents evaluate long-term rewards (R_L) through the β-δ discounting mechanism, where present bias and time discounting jointly determine utility.

The agent's utility under a subjective and evolving threshold can be expressed as:

$$U_i = \begin{cases} U(R_S), & z > \theta_t \\ \beta\delta^t U(R_L), & z \leq \theta_t \end{cases}$$

This representation captures a discontinuity in preferences, where even marginal changes in uncertainty near the threshold can trigger abrupt shifts in decision-making behavior. Unlike static models where the threshold remains fixed, the framework considers the threshold (θ_t) as a dynamic variable that evolves over time based on past experiences, current uncertainty levels, and external shocks.

The evolution of the threshold is formalized as:

$$\theta_{t+1} = \alpha\theta_t + (1-\alpha)z_t + (1-\alpha)\epsilon_t + \eta_t \tag{1}$$

In this recursive relationship, α represents the weight assigned to the previous threshold (θ_t), reflecting the agent's inertia or resistance to changing their internal benchmark. The term z_t captures the current level of perceived uncertainty, serving as an immediate adjustment factor. Here, a noise term (ϵ_t) is also introduced to captures random fluctuations in subjective uncertainty perception around z_t, weighted by $(1-\alpha)$. The shock term (η_t) introduces exogenous variability, accounting for unexpected changes or external disruptions that influence the subjective perception of uncertainty. This recursive structure highlights the path dependence of the threshold, where historical uncertainty realizations and shocks influence future decision boundaries. To better contextualize the dynamic behavior of the subjective threshold, the derivation of its long term mean is provided in Appendix A.

To facilitate analytical tractability, several assumptions are imposed on the stochastic terms within the model. Firstly, the noise term (ϵ_t) and the shock term (η_t) are posited to be mutually independent random variables. This assumption of independence ensures a null covariance between these terms, thereby permitting a separate analysis of their respective contributions to the variance of the subjective threshold. Secondly, the perceived uncertainty term (z_t), representing the agent's subjective assessment of environmental conditions, is assumed to be either deterministic or, alternatively, uncorrelated with both the noise term (ϵ_t) and the shock term (η_t). This condition consequently eliminates covariance terms involving z_t from the variance analysis.

These independence assumptions are essential for simplifying the recursive variance expression of the subjective threshold. They allow each stochastic component to be isolated, enabling a clear decomposition of how noise variance and external shocks contribute to the overall vari-

ability of the threshold. Without these assumptions, covariance terms would introduce additional dependencies, complicating both the theoretical analysis and the interpretation of results.

The interaction between dynamic uncertainty and an evolving threshold introduces a feedback loop into decision-making behavior. In stable environments, where uncertainty remains consistently low, the threshold gradually converges to a steady-state level, reducing the frequency of preference reversals. Conversely, in volatile environments characterized by frequent shocks, the threshold becomes more unstable, amplifying preference reversals and increasing reliance on short-term outcomes.

Dynamic optimization further refines the understanding of how agents navigate evolving uncertainty. The recursive decision-making process can be expressed as:

$$V_0 = \max\{U(R_S), \beta\delta^t \mathbb{E}[V_1(z_{t+1}, \theta_{t+1})]\}$$

In this framework, the Present Self evaluates immediate returns (R_S) against the discounted expected utility of long-term returns (R_L), while accounting for both the future evolution of uncertainty (z_{t+1}) and the dynamic adjustment of the threshold (θ_{t+1}). This recursive relationship underscores the dual role of uncertainty and subjective thresholds in shaping time preferences across different temporal horizons.

The condition for preferring long-term rewards under low uncertainty is maintained, such that:

$$R_L \geq \frac{R_S}{\beta\delta^t}$$

This condition indicates that long-term rewards must compensate not only for time discounting (δ^t) but also for present bias (β) to remain attractive under conditions of low uncertainty. When uncertainty exceeds the subjective threshold, however, agents systematically default to short-term preferences regardless of the magnitude of long-term rewards.

The recursive evolution of the threshold introduces several behavioral implications. First, preference instability arises near the threshold, where even small fluctuations in perceived uncertainty can trigger abrupt reversals in behavior. Second, path dependence implies that historical uncertainty shocks influence agents' future behavior, creating a cumulative effect on the threshold's trajectory. Third, feedback loops suggest that repeated exposure to high uncertainty can elevate the threshold over time, entrenching short-term preferences and reducing the likelihood of future long-term commitments.

The theoretical contribution of this framework lies in its integration of dynamic subjective uncertainty thresholds into the β-δ discounting model. By allowing the threshold to evolve over time, the framework accounts for both intrapersonal conflict, driven by present bias and time discounting, and the dynamic perception of uncertainty, shaped by past experiences and environmental shocks. This dynamic adjustment process captures both preference reversals caused by marginal shifts in uncertainty and structural breaks triggered by significant threshold-crossing events.

The recursive nature of the threshold adjustment also highlights the importance of stabilizing

mechanisms in reducing uncertainty-driven preference reversals. Policies aimed at minimizing volatility in perceived uncertainty or improving risk communication can prevent unnecessary threshold adjustments. Similarly, interventions designed to address behavioral biases, such as framing effects or ambiguity aversion, can stabilize agents' perceptions of uncertainty and reduce reliance on short-term rewards.

In summary, the theoretical framework formalizes a dynamic subjective threshold mechanism within the β-δ utility model, providing a refined understanding of how evolving uncertainty shapes intertemporal decision-making. This structure reveals the recursive interaction between subjective risk perceptions, historical uncertainty shocks, and present bias, offering insights into the behavioral mechanisms underlying time preference inconsistency and the conditions under which long-term commitments become fragile.

4 Theoretical Insights into Dynamic Preferences and Noise Effects

Proposition 1. *(Noise Variance and Its Asymmetric Influence on Preference Switching) When perceived uncertainty (z) approaches the subjective threshold (θ_t), the effect of increasing noise variance (σ^2) on the likelihood of switching between short-term and long-term rewards depends on the relative position of z to θ_t. If $z > \theta_t$, increasing noise variance reduces the likelihood of switching to long-term rewards. Conversely, if $z < \theta_t$, increasing noise variance increases the likelihood of switching to short-term rewards. At the exact threshold ($z = \theta_t$), noise variance has no meaningful effect on switching probabilities.*

Proof. The recursive optimization problem evaluates the trade-off between short-term and long-term rewards as:

$$V_0 = \max\left\{U(R_S), \beta\delta^t \mathbb{E}[V_1(z_{t+1}, \theta_{t+1})]\right\}$$

The subjective threshold (θ_t) separates short-term and long-term preferences. When uncertainty approaches the threshold ($z \approx \theta_t$), small fluctuations due to noise can disrupt preference consistency.

Noise variance modifies perceived uncertainty:

$$z' = z + \epsilon, \quad \epsilon \sim N(0, \sigma^2)$$

The probability of switching preferences depends on whether perceived uncertainty exceeds the subjective threshold and is expressed as:

$$P(z' > \theta_t) = 1 - \Phi\left(\frac{\theta_t - z}{\sigma}\right)$$

where $\Phi(\cdot)$ denotes the standard normal cumulative density function. Differentiating this proba-

bility with respect to noise variance (σ^2) using the chain rule yields:

$$\frac{\partial P(z' > \theta_t)}{\partial \sigma^2} = \phi\left(\frac{\theta_t - z}{\sigma}\right) \cdot \frac{\theta_t - z}{2\sigma^3}$$

In this expression, $\phi(\cdot)$ represents the standard normal probability density function, which is always positive. The term $\frac{1}{2\sigma^3}$ is also always positive, given that $\sigma > 0$. Therefore, the sign of the derivative depends entirely on the sign of the numerator, $\theta_t - z$.

When the perceived uncertainty (z) exceeds the threshold (θ_t), the difference $\theta_t - z$ becomes negative. Substituting this into the derivative reveals that the effect of increasing noise variance reduces the probability of switching preferences. In this region, noise variance reinforces the agent's preference for short-term rewards, effectively stabilizing their current choice.

Conversely, when perceived uncertainty (z) is below the threshold (θ_t), the difference $\theta_t - z$ becomes positive. Substituting into the derivative shows that increasing noise variance increases the probability of switching preferences. In this region, noise variance introduces instability into the agent's decision-making, raising the likelihood of switching away from long-term preferences.

At the exact threshold, where $z = \theta_t$, the difference $\theta_t - z$ becomes zero. Substituting into the derivative causes the entire term to collapse to zero. At this point, increasing noise variance has no meaningful effect on the likelihood of switching preferences.

This analysis highlights the nuanced relationship between noise variance and switching probabilities. Above the threshold, noise variance stabilizes short-term preferences, while below the threshold, it destabilizes long-term preferences. At the threshold itself, noise variance becomes economically insignificant and does not meaningfully affect preference dynamics. □

The relationship between noise variance and switching probability exhibits distinct behaviors depending on the relative position of perceived uncertainty (z) to the subjective threshold (θ_t). When perceived uncertainty exceeds the threshold ($z > \theta_t$), increasing noise variance reduces the likelihood of switching preferences, reinforcing the dominance of short-term rewards. In contrast, when perceived uncertainty remains below the threshold ($z < \theta_t$), increasing noise variance raises the likelihood of switching preferences, introducing instability into long-term commitments. At the exact threshold ($z = \theta_t$), noise variance ceases to have any meaningful impact on switching probabilities, as the derivative becomes zero. This result highlights the dominance of the threshold's position over noise variance in determining preference stability. Noise variance primarily acts as an amplifier or stabilizer depending on whether perceived uncertainty is above or below the threshold, but at the threshold itself, its influence dissipates entirely.

This initial analysis reveals the nuanced and asymmetric ways in which noise can influence an individual's preference switching at the subjective threshold. With a clear understanding of the immediate effects of noise, the following proposition explores how these effects translate to observable behavioral instability.

Proposition 2. *(Behavioral Instability Near the Threshold) Agents exhibit increased behavioral instability*

near the subjective uncertainty threshold (θ_t), and this instability grows disproportionately with higher noise variance (σ^2).

Proof. The agent's preference for long or short term rewards is determined by whether perceived uncertainty (z') falls above or below the subjective threshold (θ_t). Specifically, when $z' \leq \theta_t$, the agent prefers the long term reward; otherwise they prefer the short term reward. The probability of choosing a short-term reward is the probability of switching preferences and is expressed as:

$$P(\text{Choose Short Term}) = P(z' > \theta_t) = 1 - \Phi\left(\frac{\theta_t - z}{\sigma}\right)$$

Where Φ is the CDF of the standard normal distribution, θ_t is the subjective threshold, z is the perceived uncertainty, and σ is the noise variance.

This binary decision-making process can be modeled as a Bernoulli random variable, where the outcome (U) is defined as:

$$U = \begin{cases} 1, & \text{if } z' > \theta_t \\ 0, & \text{if } z' \leq \theta_t \end{cases}$$

where $U = 1$ represents switching preferences and $U = 0$ represents no switch.

The variance of a Bernoulli random variable is given by:

$$\text{Var}(U) = P(1 - P)$$

This expression provides a simple and direct representation of how the switching probability impacts the variability of the decision outcome.

The Delta Method, introduced by Jacob Wolfowitz (1954), provides a way to approximate the variance of a function of a random variable. Specifically, the Delta Method can be used to analyze how changes in the value of z' lead to changes in the probability P.

The Delta Method states that for a differentiable function $g(x)$ of a random variable x, the variance of the function can be approximated by:

$$\text{Var}(g(x)) \approx \left(\frac{dg}{dx}\right)^2 \text{Var}(x)$$

This approximation is derived from the first order Taylor Expansion of $g(x)$ around its mean, where any higher order terms are assumed to be negligible given that changes in the value of x are small. This assumes that the function $g(x)$ is reasonably well behaved and differentiable around the mean of x. Specifically, given that $g(x)$ is the switching probability function P, the first order Taylor Expansion is given by:

$$P(x) = P(\mu) + P'(\mu)(x - \mu)$$

Where μ is the mean of x, and $P'(\mu)$ is the derivative of the probability with respect to x evaluated

at the mean. As such, this shows that the variance is approximately given by:

$$\text{Var}(P(x)) \approx (P'(\mu))^2 \text{Var}(x)$$

In this context, $g(x)$ corresponds to the switching probability (P), and x corresponds to the perceived uncertainty (z'), influenced by noise variance (σ^2). The variance of the switching probability can be approximated as:

$$\text{Var}(P) \approx \left(\frac{\partial P}{\partial z}\right)^2 \text{Var}(z')$$

Given that the perceived uncertainty z' is dominated by the noise term ϵ, its variance is approximated by the variance of noise:

$$\text{Var}(z') \approx \text{Var}(\epsilon) = \sigma^2$$

Substituting this into the variance expression gives:

$$\text{Var}(P) \propto \left(\frac{\partial P}{\partial z}\right)^2 \sigma^2$$

where $\propto$ means proportional to.

Differentiating the switching probability with respect to perceived uncertainty yields:

$$\frac{\partial P}{\partial z} = \frac{\partial}{\partial z}\left(1 - \Phi\left(\frac{\theta_t - z}{\sigma}\right)\right) = \phi\left(\frac{\theta_t - z}{\sigma}\right) \cdot \frac{1}{\sigma}$$

where ϕ is the PDF of the normal distribution. Substituting this derivative into the variance formula, the variance becomes:

$$\text{Var}(P) \propto \phi^2\left(\frac{\theta_t - z}{\sigma}\right)$$

The standard normal probability density function (ϕ) peaks near $z \approx \theta_t$, where small changes in perceived uncertainty have a disproportionately large effect on switching probability. In this region, the variance increases due to the large value of the derivative $\frac{\partial P}{\partial z}$.

Graphically, the probability density function has a bell curve shape, where the peak of the curve represents a region of maximal sensitivity to changes in the argument. With respect to the probability of a switch, this curve determines how sensitive the result is to changes in the value of z.

As noise variance (σ^2) increases, the variance of the switching probability becomes proportionally larger, since the term σ^2 is outside the PDF, while the peak height of the PDF itself is reduced, as ϕ is inversely proportional to σ. These effects combine to mean that changes near the threshold can have an out-sized impact on probability at high variance. Therefore, behavioral instability scales non-linearly due to the combined effects of the heightened sensitivity near the threshold and the direct effect of variance. □

This proposition demonstrates that when individuals are uncertain about their preferences, their decisions can become highly susceptible to even small amounts of random noise. Specifically, when perceived uncertainty is close to their subjective threshold, decision making becomes erratic and unpredictable. This effect is amplified at higher levels of noise variance, making preferences less stable and more prone to sudden fluctuations. This model shows how uncertainty undermines clear preferences and leads to inconsistent choices, particularly when they are not sure about what the best choice may be.

Having shown that fluctuations are greatest near the threshold, the subsequent proposition investigates how these short term disturbances, as well as an agent's own past experience, accumulate over time and create long term effects on the overall decision making process, and the stability of their internal benchmark.

Proposition 3. *(Preference Instability Under Variable Noise) As noise variance (σ^2) increases from low levels towards infinity, agents exhibit a shift from deterministic preferences based on the relationship between the subjective threshold (θ_t) and perceived uncertainty (z) towards increasingly random behavior. For large but finite noise variance ($\sigma \gg 1$), deviations from random choice are proportional to $\frac{\theta_t - z}{\sigma}$, and as noise variance approaches infinity ($\sigma \to \infty$), the probability of choosing either long term or short term rewards, and of exhibiting preference reversals, converges to precisely 50%, independent of discount factors (δ), present bias (β), or the relationship between perceived uncertainty and the threshold.*

Proof. The agent's preference for long or short term rewards is determined by whether perceived uncertainty (z') falls above or below the subjective threshold (θ_t). Specifically, when $z' \leq \theta_t$, the agent prefers the long term reward; otherwise they prefer the short term reward. The probability of choosing a long term reward is:

$$P(\text{Choose Long Term}) = \Phi\left(\frac{\theta_t - z}{\sigma}\right)$$

Where Φ is the CDF of the standard normal distribution, θ_t is the subjective threshold, z is the perceived uncertainty, and σ is the noise variance. This represents the probability of choosing a long-term reward, and therefore the probability of not switching preference. The probability of choosing a short term reward (and thus of switching preference) is:

$$P(\text{Choose Short Term}) = 1 - \Phi\left(\frac{\theta_t - z}{\sigma}\right)$$

To understand how noise variance influences both preference expression and reversals, it is necessary to examine the behavior of these probabilities under varying levels of noise.

When noise variance is low ($\sigma \approx 0$), then preference choices are highly predictable. As σ approaches zero, the probability of the long term reward is:

$$\lim_{\sigma \to 0} P(\text{Choose Long Term}) = \begin{cases} 1 & \text{if } z < \theta_t \\ 0 & \text{if } z > \theta_t \end{cases}$$

and for the short term reward is:

$$\lim_{\sigma \to 0} P(\text{Choose Short Term}) = \begin{cases} 1 & \text{if } z > \theta_t \\ 0 & \text{if } z < \theta_t \end{cases}$$

This demonstrates that at low noise, choices are completely determined by the relationship between perceived uncertainty and the subjective threshold. If the perceived uncertainty is greater than the threshold, short term will always be chosen, and vice versa. Preference reversals are highly unlikely.

To evaluate the behavior under high noise variance, the Taylor series expansion can be used from the previous proposition:

$$\Phi\left(\frac{\theta_t - z}{\sigma}\right) \approx \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \frac{\theta_t - z}{\sigma} + O\left(\left(\frac{\theta_t - z}{\sigma}\right)^3\right)$$

For large but finite noise variance ($\sigma \gg 1$), the term $\frac{\theta_t - z}{\sigma}$ determines deviations from randomness. As variance increases, the magnitude of this term shrinks, thus reducing the influence of the subjective threshold.

As noise variance approaches infinity ($\sigma \to \infty$), the term $\frac{\theta_t - z}{\sigma}$ approaches zero, and the CDF converges to:

$$\lim_{\sigma \to \infty} \Phi\left(\frac{\theta_t - z}{\sigma}\right) = \frac{1}{2}$$

This means that the probability of choosing long term rewards is:

$$\lim_{\sigma \to \infty} P(\text{Choose Long Term}) = \frac{1}{2}$$

and the probability of choosing short term rewards and switching preference is:

$$\lim_{\sigma \to \infty} P(\text{Choose Short Term}) = 1 - \frac{1}{2} = \frac{1}{2}$$

This demonstrates that at infinite noise variance, the probability of selecting the long-term reward, and the probability of exhibiting preference reversals both converges to 50%. Therefore, regardless of the relationship between the subjective threshold and the perceived uncertainty, choices become completely random.

The derivative with respect to σ of the switching probability can be considered:

$$\frac{\partial}{\partial \sigma}\left(1 - \Phi\left(\frac{\theta_t - z}{\sigma}\right)\right) = \phi\left(\frac{\theta_t - z}{\sigma}\right) \frac{(\theta_t - z)}{\sigma^2}$$

This derivative shows how the direction of the shift in probability is dependent on the sign of $\theta_t - z$, although at extremely high noise this has limited impact.

This combined analysis demonstrates that as noise variance increases from low levels towards

infinity, the model moves from one of deterministic choice, dependent on the subjective threshold, towards a state of completely random choice independent of present bias, discount factors, the subjective threshold, and perceived uncertainty. The result is a breakdown of stable discounting, leading to preferences that change erratically over time and undermine an agent's ability to commit to long-term plans, highlighting a key mechanism behind time-inconsistent preferences. At extremely high levels of noise, agents' decisions become entirely driven by random fluctuations, which directly subverts any attempt to produce time-consistent preferences. □

This proposition reveals that while individuals might try to make rational decisions by weighing risks and future rewards, high levels of unpredictable environmental noise can completely disrupt this process. As the level of noise goes up, decision making shifts from a system where stable plans can be made, towards an increasingly random and chaotic system, completely undermined by stochastic fluctuations. This effect makes long term planning very difficult, because decisions become volatile and unpredictable and thus undermines individuals' ability to maintain consistent plans for the future.

With the understanding that noise can affect decision making both locally and globally, the next proposition builds on the previous results to show how stochastic shocks and noise, interacting with a persistence parameter, can result in a threshold that is unstable in the long run.

Proposition 4. *(Persistence, Stochasticity, and the Fragility of Long-Term Commitments) Historical uncertainty shocks (η_t) and noise variance (σ^2) exert a persistent influence on the subjective threshold (θ_t). While agents can anticipate the average trajectory of the threshold, stochastic fluctuations introduced by noise and shocks amplify variability and induce upward drift in the threshold, preventing convergence to a stable benchmark. This persistent variability entrenches short-term preferences and reduces long-term commitments, fostering fragility in intertemporal decision-making.*

Proof. As established in the theoretical framework (Equation (1)), the threshold evolves according to:

$$\theta_{t+1} = \alpha\theta_t + (1-\alpha)z_t + (1-\alpha)\epsilon_t + \eta_t$$

where $\alpha \in [0,1]$ represents the degree of inertia in adjusting the threshold, z_t denotes perceived uncertainty at time t, ϵ_t represents a noise term, and η_t captures random exogenous shocks. The noise term modifies perceived uncertainty as follows: $\epsilon_t \sim N(0, \sigma^2)$. The random variables ϵ_t and η_t are assumed to be independent. Taking the expectation of both sides yields:

$$\mathbb{E}[\theta_{t+1}] = \alpha\mathbb{E}[\theta_t] + (1-\alpha)\mathbb{E}[z_t] + (1-\alpha)\mathbb{E}[\epsilon_t] + \mathbb{E}[\eta_t]$$

Since $\mathbb{E}[\epsilon_t] = 0$ and $\mathbb{E}[\eta_t] = 0$:

$$\mathbb{E}[\theta_{t+1}] = \alpha\mathbb{E}[\theta_t] + (1-\alpha)\mathbb{E}[z_t]$$

This equation shows that the expected value of the threshold is determined by a weighted average of the previous expected threshold and perceived uncertainty. Agents can use this to anticipate

the mean trajectory of θ_t. However, the stochastic components ϵ_t and η_t introduce variability that is not captured by the mean, preventing precise prediction of the threshold at any specific time.

To analyze the variability, consider the variance of the threshold at time $t+1$. Assuming ϵ_t and η_t are independent, the variance of the threshold at $t+1$ can be expressed as:

$$\text{Var}(\theta_{t+1}) = \alpha^2 \text{Var}(\theta_t) + (1-\alpha)^2 \text{Var}(\epsilon_t) + \text{Var}(\eta_t)$$

Substituting the variance of ϵ_t as σ^2:

$$\text{Var}(\theta_{t+1}) = \alpha^2 \text{Var}(\theta_t) + (1-\alpha)^2 \sigma^2 + \text{Var}(\eta_t)$$

Expanding this expression recursively to n periods yields:

$$\text{Var}(\theta_{t+n}) = \alpha^{2n} \text{Var}(\theta_t) + (1-\alpha)^2 \sigma^2 \sum_{i=0}^{n-1} \alpha^{2i} + \sum_{i=0}^{n-1} \alpha^{2i} \text{Var}(\eta_{t+i})$$

If the shock variance remains constant over time, denoted as $\text{Var}(\eta_{t+i}) = \text{Var}(\eta)$, the variance of the threshold at time $t+n$ can be simplified and expressed as:

$$\text{Var}(\theta_{t+n}) = \alpha^{2n} \text{Var}(\theta_t) + (1-\alpha)^2 \sigma^2 \frac{1-\alpha^{2n}}{1-\alpha^2} + \text{Var}(\eta) \frac{1-\alpha^{2n}}{1-\alpha^2}$$

Note that the term α^{2n} decreases over time as n increases and $\alpha \in [0,1]$. The term $\alpha^{2n} \text{Var}(\theta_t)$ represents how much of the historical variance from the initial period t is carried forward. This term highlights how past noise and shocks persist in influencing the threshold over time. As α approaches 1, the rate at which this initial variance dissipates decreases, increasing the cumulative impact of new noise and shocks on the overall variance. Thus, the term does not represent a new source of stochastic variance at $t+n$, but rather the persistence of past variance.

The instantaneous variance growth based on the previous time period is given by:

$$\text{Var}(\theta_{t+1}) = (1-\alpha)^2 \sigma^2 + \text{Var}(\eta_t)$$

This demonstrates that the threshold variance is directly influenced by the variance of noise and external shocks. Differentiating the threshold's variance with respect to the noise variance reveals:

$$\frac{\partial \text{Var}(\theta_{t+1})}{\partial \sigma^2} = (1-\alpha)^2$$

Since $(1-\alpha)^2 > 0$ for valid α in $[0,1]$, an increase in noise variance directly contributes to increased variability in the threshold's trajectory. The long-term variance can then be considered a scaled version of this instantaneous variance.

Considering the variance of the threshold after a long period of time ($n \to \infty$):

$$\lim_{n\to\infty} \text{Var}(\theta_{t+n}) = \frac{(1-\alpha)^2\sigma^2 + \text{Var}(\eta)}{1-\alpha^2}$$

This equation demonstrates how the persistence of the threshold (captured by α) acts as an amplification factor, increasing the influence of noise (σ^2) and uncertainty shocks ($\text{Var}(\eta)$) on long-run variability. This reinforces the path-dependent nature of the threshold and the lack of convergence to a stable state, as increases to the persistence parameter also increases the sensitivity of the threshold to shocks and noise.

To understand how stochastic terms contribute to an upward drift in the threshold, the accumulated impact of shocks can be examined. From Equation (1), the recursive expansion of the threshold can be expressed as:

$$\theta_{t+n} = \alpha^n\theta_t + \sum_{i=0}^{n-1} \alpha^i(1-\alpha)z_{t+i} + \sum_{i=0}^{n-1} \alpha^i(1-\alpha)\epsilon_{t+i} + \sum_{i=0}^{n-1} \alpha^i\eta_{t+i}$$

This equation shows how each prior shock impacts the current threshold, where the effect of each shock is scaled by the persistence parameter α over time. The accumulation of these stochastic values is directly reflected in the threshold's value over time. Although noise terms and shocks have a mean of zero, the positive values are not dissipated, and will contribute to the long term upward drift of the threshold.

In the long term, the combination of positive stochastic terms and the persistence parameter result in an increase in the mean threshold value over time, demonstrating upward drift. The recursive nature of the threshold, combined with these sources of variability, prevents convergence to a stable long-term benchmark. This, in turn, results in agents remaining entrenched in short-term decision patterns, increasing their likelihood of deviating from sustained long-term commitments. The persistence parameter α is crucial in determining the extent to which past shocks and noise drive this process. The mechanisms explored in this proposition highlight the ways in which cognitive biases, the accumulation of stochastic shocks, and the reliance on a fluctuating subjective threshold can interfere with an agent's ability to make consistent decisions, leading to deviations from long-term plans by introducing time-inconsistent preferences. □

This proposition suggests that when people make decisions over time, their subjective views about risk and uncertainty are not stable. Instead, they are continuously influenced by past experiences and random shocks in the environment. This influence is amplified by the fact that current decisions are strongly tied to past ones and means that they can become stuck with a preference set that continually pushes them in the same direction, even if it's not optimal for the long term. This makes it difficult to maintain long-term commitments because their decision-making processes become sensitive to short-term fluctuations in perceived risk, rather than a consistent approach. This recursive accumulation prevents the subjective threshold from stabilizing, leading to entrenched short-term preferences and reduced flexibility in long-term commitments. The persistence param-

eter (α) plays a crucial role in determining whether agents can overcome noise-induced deviations or remain locked into recursive patterns of suboptimal decision-making. This means that an individual's subjective discount rate, driven by this threshold, is not stable, leading to inconsistent decisions over time, where they prioritize short term goals and neglect long term ones.

The previous propositions have explored the dynamics of the threshold and their local impact on decision making. Given the unstable nature of this decision making mechanism, the subsequent proposition investigates how an agent can improve the stability of the threshold by adapting to past values.

Proposition 5. *(Adaptive Persistence and Subjective Threshold Stability) When agents adjust their persistence parameter (α) in response to past subjective thresholds (θ_t), long-term threshold stability improves. Specifically, an adaptive α that decreases after prolonged periods of high or low thresholds reduces the long-term variance of the subjective threshold (θ_t), when compared to a model with a fixed persistence parameter.*

Proof. As established in the theoretical framework (Equation (1)), the threshold evolves according to the recursive relationship.

To model adaptive persistence, let the value of α_t be a function of past threshold values such that:

$$\alpha_t = \alpha_0 - \gamma \left(\theta_t - \theta_0\right)^2$$

where $\alpha_0 \in [0, 1]$ is a base level of persistence, $\gamma > 0$ is a scaling parameter, and θ_0 represents an optimal level of the threshold. This equation, the derivation of which is provided in Appendix B, indicates that if the threshold has recently deviated significantly from the optimal level, the value of α_t will be reduced, thus lowering the degree of persistence. Conversely, high levels of persistence will be associated with a subjective threshold that is relatively stable around the long term value of θ_0.

The variance of the threshold at time $t+1$ can be expressed as:

$$\text{Var}(\theta_{t+1}) = \alpha_t^2 \text{Var}(\theta_t) + (1-\alpha_t)^2 \text{Var}(z_t) + (1-\alpha_t)^2 \sigma^2 + \text{Var}(\eta_t)$$

Assuming that the shock variance remains constant over time and perceived uncertainty is influenced by a stochastic noise term, where $\epsilon_t \sim N(0, \sigma^2)$, the relationship can be expressed as:

$$\text{Var}(\theta_{t+1}) = \alpha_t^2 \text{Var}(\theta_t) + (1-\alpha_t)^2 \sigma^2 + \text{Var}(\eta_t)$$

As shown in the previous proposition, the long-term variance of the threshold, with a fixed parameter of persistence, is given by:

$$\lim_{n \to \infty} \text{Var}(\theta_{t+n}) = \frac{(1-\alpha)^2 \sigma^2 + \text{Var}(\eta)}{1-\alpha^2}$$

This expression highlights that high values of α are associated with higher variance in the long term, as they amplify the effects of stochastic variance introduced each period.

In contrast to the model with a fixed level of persistence, the adaptive persistence model, results in a lower value of α_t at times when the threshold is highly variable, and a higher value when it is more stable. With this adaptive mechanism, high variance in the threshold will result in reduced persistence, and the influence of past values becomes less important over time.

To better understand the impact of the adaptive persistence, the long term threshold value can be analyzed using a recursive argument. By recursively inserting the value of θ_{t+n} into the expression for θ_{t+n+1}, the long-term threshold value can be expressed as:

$$\theta_{t+n} = \left(\prod_{i=0}^{n-1} \alpha_{t+i}\right) \theta_t + \sum_{i=0}^{n-1} \left(\left(\prod_{j=i+1}^{n-1} \alpha_{t+j}\right) (1-\alpha_{t+i}) z_{t+i} \right) + \sum_{i=0}^{n-1} \left(\left(\prod_{j=i+1}^{n-1} \alpha_{t+j}\right) (1-\alpha_{t+i}) \epsilon_{t+i} + \left(\prod_{j=i+1}^{n-1} \alpha_{t+j}\right) \eta_{t+i} \right)$$

Where $\prod_{j=i+1}^{n-1} \alpha_{t+j} = 1$ when $i = n-1$. This equation highlights how a time-varying value of α_t will impact future threshold values. Specifically, if it is assumed that the threshold will tend towards a mean value, then it can be assumed that when there is high variance, the values of alpha will be lower. This means that the product of α_t terms will decrease in response to high variance, and the influence of past thresholds, as represented by θ_t, will be reduced.

Therefore, while a fixed level of persistence amplifies variance, as demonstrated in Proposition 4, this adaptive model reduces this amplification effect, and the long term variance. The negative relationship between the level of variance and the persistence parameter implies that a model with adaptive persistence will tend to exhibit lower long term variability.

In summary, the recursive model that includes adaptive persistence, where the parameter α_t varies as a function of the subjective threshold, is more stable than the model with a fixed level of persistence. The model with adaptive persistence exhibits reduced long-term variability due to the negative relationship between the level of variance and the persistence parameter. The model with an adaptive persistence parameter ensures that past shocks and noise have a smaller impact on the long term threshold. □

This proposition reveals a mechanism by which agents can counteract the destabilizing effects of past uncertainty. Stable decision-making can be achieved not by blindly adhering to past experiences, but by actively adapting to current volatility. By reducing their reliance on past threshold values, and specifically by reducing the persistence parameter when their thresholds are volatile, agents limit the propagation of these fluctuations, thus achieving greater long-term stability. This active adjustment of the persistence parameter helps to insulate the decision-making process from the accumulation of past stochasticity. They can mitigate the impact of short-term shocks on their long-term planning, allowing for a more future-oriented decision-making approach, and in so doing, can increase preference consistency. This also leads to a more stable threshold, which can promote more consistent behavior.

In showing the utility of adaptive persistence, the subsequent proposition investigates the

dynamics of how the threshold will return to a more consistent long term state after a shock, and how that is influenced by a persistence parameter.

Proposition 6. *(Persistence and the Speed of Subjective Threshold Recovery) The speed at which the subjective threshold (θ_t) returns to its expected value after a shock is inversely related to the persistence parameter (α). Specifically, a lower persistence parameter results in faster convergence of the threshold to its expected value and faster recovery to a baseline level, while higher persistence parameter creates slower recovery and longer periods of deviation after a shock.*

Proof. As established in the theoretical framework (Equation (1)), the threshold evolves according to the recursive relationship.

To analyze the speed of recovery after a shock, it can be assumed that at time $t = 0$ there is a single shock of size η_0, a noise term of ϵ_0 and the threshold is at some starting value θ_0, such that:

$$\theta_1 = \alpha\theta_0 + (1 - \alpha)\epsilon_0 + \eta_0$$

After two periods, the threshold will be:

$$\theta_2 = \alpha\theta_1 = \alpha^2\theta_0 + \alpha(1 - \alpha)\epsilon_0 + \alpha\eta_0$$

Recursively substituting this value at n periods:

$$\theta_n = \alpha^n\theta_0 + (1 - \alpha)\sum_{i=0}^{n-1}\alpha^i\epsilon_0 + \sum_{i=0}^{n-1}\alpha^i\eta_0$$

This equation shows how a single shock at time zero impacts the value of the threshold in subsequent time periods. Specifically, the persistence parameter α directly scales the impact of the initial threshold and the initial shock at each successive period. As the number of time periods n increases, the term α^n approaches zero given that $|\alpha| < 1$, which indicates that the impact of the starting threshold on the present value decreases over time. In a similar manner, the magnitude of prior shocks is decreased in every subsequent period.

A recovery time can be defined as the time at which the effect of the shock becomes less than some arbitrary small value.

To show more clearly the effect of the persistence parameter, it is helpful to define a recovery time based on the long-term expected value of the threshold in the absence of any shocks. The long term expected value of the threshold is such that:

$$E[\theta_{t+1}] = \alpha E[\theta_t] + (1 - \alpha)E[z_t]$$

If it is assumed that the long term value of the perceived uncertainty is zero, then this expression can be simplified to:

$$E[\theta_{t+1}] = \alpha E[\theta_t]$$

And if the long term value is considered, it can be shown that the expected threshold will approach zero given that $|\alpha| < 1$:

$$\lim_{n\to\infty} E[\theta_{t+n}] = 0$$

Therefore, a recovery time can be defined as the number of time periods that it takes for the expected value of the threshold to return to zero, based on the initial values of the threshold and the initial shock.

To analyze this value more formally, it can be considered that at some time n the impact of the initial threshold is approximately zero, at which point the equation can be simplified to:

$$\theta_n \approx (1-\alpha)\sum_{i=0}^{n-1}\alpha^i\epsilon_0 + \sum_{i=0}^{n-1}\alpha^i\eta_0$$

If it is assumed that the initial noise shock was negligible, leaving only the shock term, the expression can be further simplified as:

$$\theta_n \approx \sum_{i=0}^{n-1}\alpha^i\eta_0$$

The threshold can be said to have approximately "recovered" when this remaining sum of past shocks is less than some very small value, say ϵ, such that:

$$\sum_{i=0}^{n-1}\alpha^i\eta_0 < \epsilon$$

This is equivalent to:

$$\begin{aligned} & \eta_0\frac{1-\alpha^n}{1-\alpha} < \epsilon, \\ \implies \quad & \alpha^n > 1 - \frac{\epsilon(1-\alpha)}{\eta_0}, \\ \implies \quad & n > \frac{\ln\left(1-\frac{\epsilon(1-\alpha)}{\eta_0}\right)}{\ln(\alpha)}. \end{aligned}$$

Using L'Hôpital's rule:

$$\lim_{\alpha\to 1}\frac{\ln\left(1-\frac{\epsilon(1-\alpha)}{\eta_0}\right)}{\ln(\alpha)} = \lim_{\alpha\to 1}\frac{\frac{\epsilon}{\eta_0}}{1-\frac{\epsilon(1-\alpha)}{\eta_0}}\frac{\alpha}{1} = \frac{\frac{\epsilon}{\eta_0}}{1}\frac{1}{1} = \frac{\epsilon}{\eta_0}$$

Thus, as α approaches 1, the value of n approaches infinity.

Alternatively, a similar result can be derived by calculating the half-life of the shock, which is defined as the number of time periods that is needed before the effect of the shock is reduced to

one half of its original impact:

$$\alpha^n = \frac{1}{2}$$

$$n = \frac{ln(1/2)}{ln(\alpha)}$$

This formula also shows the inverse relationship between the value of α and the speed of recovery.

Therefore, a higher persistence parameter results in a longer recovery time and slower return to the baseline level, and a lower persistence parameter results in a faster recovery time. □

This proposition demonstrates that the speed at which agents revert to their baseline subjective uncertainty after a shock is directly influenced by their degree of persistence. Agents with a lower persistence in their internal benchmarks will experience a more rapid return to their long-term average, while higher persistence creates longer lasting effects of temporary shocks, slowing their return to a baseline level. This means that agents with higher levels of persistence may find themselves deviating from a long term plan after a shock for a longer period of time, while agents with lower persistence will be able to resume normal behavior faster. In essence, the persistence parameter not only impacts the stability of the threshold but also the speed at which the agent can adjust their behavior. However, the speed at which these effects can be counteracted is also determined by persistence, highlighting a key trade-off in decision making: while a higher persistence parameter can help to reduce overall variance in the threshold and improve stability, it also acts to prolong the effect of any negative shock, and increases the time required to return to baseline preferences. Conversely, a lower persistence parameter leads to greater responsiveness and faster recovery after a shock, but with a more volatile threshold which may make it harder to commit to long term plans in the first place. This trade-off implies that the optimal level of persistence for an agent will depend on the relative costs of the two behaviors, which is something that will likely be context specific.

5 Theoretical Positioning of the Dynamic Threshold Model

The dynamic threshold model offers a novel theoretical framework for understanding time preferences by introducing a dynamic subjective threshold that evolves over time in response to stochastic shocks, perceived uncertainty, and a persistence mechanism. Unlike traditional β-δ discounting models, which rely on fixed parameters for present bias and discount rates, the dynamic threshold introduces a recursive structure where past values, external disturbances, and noise continuously reshape the decision-making process. This dynamic structure captures behavioral instability near the threshold, where preference reversals are most likely to occur, and emphasizes how stochastic noise amplifies these effects non-linearly. Such an approach shifts focus away from static assumptions, highlighting the temporal and environmental dependence of preferences.

The integration of stochastic noise and dynamic thresholds into time preference modeling reveals a critical departure from static approaches. In traditional frameworks, preference reversals

are typically attributed to fixed psychological biases, such as present bias or hyperbolic discounting. However, the dynamic threshold model illustrates that reversals can emerge naturally from stochastic fluctuations in perceived uncertainty, even in the absence of intrinsic psychological biases. This perspective broadens the analytical scope of time preference theory, providing a foundation for understanding how real-world volatility and path-dependent processes interact with individual decision-making. By incorporating these recursive mechanisms, the dynamic threshold model captures both the local and global effects of uncertainty, offering a versatile tool for explaining observed inconsistencies in intertemporal choices.

5.1 Short-Term Dynamics and Instability

Proposition 1 establishes that the probability of switching preferences depends on the relative position of perceived uncertainty and the dynamic threshold, rather than being constant. While Read (2001) empirically demonstrates sub-additive and variable discounting patterns over time, the dynamic threshold model provides a theoretical foundation for these observations by showing how stochastic noise and subjective uncertainty thresholds drive such variability. In contrast to static functional assumptions, this model explains the dynamic relationship between uncertainty and preference switching, offering a mechanism that aligns with empirical findings.

Proposition 2 extends this insight by showing that behavioral instability peaks near the threshold, with sensitivity increasing as a non-linear function of stochasticity. This result resonates with O'Donoghue and Rabin (2001), who identify procrastination as a manifestation of time-inconsistent preferences driven by present bias. While O'Donoghue and Rabin emphasize an internally fixed psychological bias, Proposition 2 highlights that instability arises from external environmental volatility rather than static psychological parameters. This distinction underscores the model's unique contribution in capturing preference instability under stochastic uncertainty.

The dynamic threshold model further emphasizes that instability near the threshold is not merely a theoretical artifact but a predictable outcome of the interaction between uncertainty and adaptive thresholds. As stochastic noise intensifies, decision boundaries become blurred, amplifying the probability of abrupt shifts between short-term and long-term preferences. This dynamic sensitivity explains why individuals exhibit higher behavioral volatility in uncertain or rapidly changing environments. It also aligns with empirical findings showing that decision-makers are more prone to inconsistencies when operating under fluctuating external conditions. This theoretical insight suggests that managing perceived uncertainty through interventions such as improved risk communication or environmental stability may reduce short-term behavioral reversals.

5.2 Long-Term Dynamics and Path Dependence

Moving towards long-term dynamics, Proposition 3 examines preference instability under variable noise, illustrating how elevated stochasticity increases randomness in decision-making, ultimately undermining long-term planning. This result shares conceptual ground with Gabaix

(2019), who attributes behavioral inconsistency to limited cognitive attention. While Gabaix emphasizes how agents fail to process all relevant information under cognitive constraints, Proposition 3 highlights that even fully attentive agents can exhibit instability when exposed to heightened stochastic noise and environmental randomness. This comparison highlights two distinct mechanisms for behavioral inconsistency: internal cognitive biases versus external stochastic forces.

Proposition 4 builds on this by demonstrating how past shocks accumulate over time, shaping long-term preference trajectories and introducing path dependence. This mechanism mirrors findings in Harris and Laibson (2001), where quasi-hyperbolic discounting generates inconsistencies due to the tension between short-term and long-term objectives. However, while Harris and Laibson rely on a deterministic and static functional form for preferences, Proposition 4 introduces a dynamic stochastic threshold mechanism, where accumulated noise and shocks recursively influence long-term preferences. This extension provides a more flexible and adaptive explanation for long-term preference fragility.

In comparison to O'Donoghue and Rabin (2001), who propose that time-inconsistent behavior stems from present bias, the dynamic threshold model offers an alternative narrative. Rather than attributing inconsistency to internal psychological biases as fixed parameters, Propositions 3 and 4 demonstrate how external stochastic shocks and accumulated noise create long-term instability. This distinction highlights two fundamentally different approaches: internal psychological mechanisms versus dynamic external adaptation.

The long-term consequences of these dynamics underscore the importance of historical uncertainty in shaping present decision-making processes. As stochastic shocks accumulate, agents may become locked into behavioral patterns dominated by short-term preferences, even when long-term rewards are objectively superior. This phenomenon aligns with empirical evidence suggesting that individuals exposed to prolonged uncertainty or repeated shocks often exhibit reduced patience and increased sensitivity to immediate outcomes. The recursive nature of the dynamic threshold ensures that past experiences are not easily discounted, reinforcing the persistence of preference instability over time. Recognizing these long-term dynamics offers critical insights into policy interventions aimed at stabilizing decision environments to promote consistent long-term planning.

5.3 Adaptation and Trade-Offs

Adaptation to environmental volatility is explored in Proposition 5, which identifies the persistence parameter (α) as a critical mechanism for managing preference instability. Agents capable of strategically adjusting their persistence parameter demonstrate resilience against external shocks. Ericson and Laibson (2019) emphasize the importance of strategic adjustments in intertemporal choice models, primarily focusing on discounting behavior. Proposition 5 extends this perspective by situating adaptive persistence within a recursive stochastic structure, demonstrating how agents dynamically recalibrate their sensitivity to volatility based on environmental feedback. This mechanism provides a nuanced understanding of how persistence shapes long-term behav-

ioral stability.

Proposition 6 introduces a trade-off between stability and responsiveness, showing that while a stable threshold enhances predictability, it reduces sensitivity to external changes. This dynamic tension aligns with insights from Halevy (2015), who provides axioms for time consistency as benchmarks for rational choice. Halevy's framework assumes fixed parameters governing discounting behavior and lacks explicit treatment of stochastic uncertainty. Proposition 6 departs from these assumptions, showing that time-inconsistent behavior emerges naturally in stochastic environments, where adaptive responses balance threshold stability against responsiveness. This adjustment introduces a dynamic understanding of how agents navigate uncertainty while attempting to maintain long-term goals.

The empirical findings by Read (2001) further contextualize this trade-off, demonstrating that preference structures exhibit significant variability and are influenced by factors beyond fixed discount functions. While Read empirically identifies non-linear patterns in discounting, Propositions 5 and 6 provide a theoretical explanation for these observations. The dynamic threshold model shows that adaptive adjustments to the persistence parameter, shaped by past shocks and environmental volatility, drive these non-linear patterns. This insight provides a robust theoretical foundation for interpreting Read's empirical findings through a stochastic and adaptive lens.

The trade-off between stability and responsiveness highlights a fundamental tension in dynamic decision environments. While higher persistence levels (α) can buffer agents from short-term volatility and reduce the frequency of abrupt preference reversals, they simultaneously slow down the adjustment process in response to significant environmental changes. Conversely, lower persistence allows for rapid adaptation but increases behavioral volatility, making long-term planning more fragile. These opposing dynamics suggest that optimal persistence levels are context-dependent, influenced by factors such as the frequency of shocks, the magnitude of uncertainty, and the agent's ability to anticipate future states. Understanding this trade-off provides valuable insights into designing behavioral interventions or decision architectures that strike an effective balance between stability and adaptability.

5.4 Comparison with Laibson and Maxted (2023)

While the dynamic threshold model builds upon several key concepts from the existing literature, including those of time-inconsistent preferences and dynamic optimization, recent research such as that of Laibson and Maxted (2023) has demonstrated the importance of time-step specification when numerically modeling present bias. They focused on the methodological problem of producing robust numerical simulations, whereas the dynamic threshold model in this paper is primarily concerned with the psychological and behavioral processes that underpin dynamic preference formation, and the theoretical basis for the internal benchmarks that may be used by decision makers. Specifically, their work highlights the sensitivity of numerical results to small changes in the specification of model parameters, and how a correct choice of parameters can produce more robust results. The dynamic threshold model, in contrast, is primarily concerned

with the mechanisms that give rise to preferences that are not robust in the first place, and seeks to demonstrate how stochastic shocks and past choices affect the internal parameters of dynamic models. By focusing on the underlying cognitive process, and explicitly modeling the accumulation of stochastic forces over time, the dynamic threshold model adds significant explanatory power to the previous methodological insights offered by Laibson and Maxted (2023). The dynamic threshold model can also explain the findings of that paper by providing a clear reason for why time-steps may have an out-sized impact on the model behavior, and by implication, the preferences that are being represented. This highlights the value of the dynamic threshold model for not only capturing behavioral tendencies, but also for explaining their numerical and methodological origins.

5.5 Remarks on the Model's Contributions

Taken together, these propositions collectively position the dynamic threshold model as a distinctive contribution to time preference theory. By introducing stochastic shocks, adaptive persistence, and a recursive subjective threshold mechanism, the model provides a comprehensive explanation for both short-term preference reversals and long-term instability. In contrast to static frameworks, such as those proposed by Harris and Laibson (2001), O'Donoghue and Rabin (2001), and Halevy (2015), the dynamic threshold model emphasizes that preferences are not fixed but context-sensitive, shaped by the interaction of stochastic noise, adaptive persistence, and environmental volatility. This approach not only addresses critical gaps in existing models but also offers a dynamic, context-aware perspective on temporal decision-making.

Furthermore, the dynamic threshold model extends beyond descriptive accuracy by offering a structural foundation for both behavioral prediction and policy intervention. The recursive nature of the subjective threshold allows for the identification of critical periods where interventions—such as stabilizing external volatility, improving agents' risk perception, or adjusting persistence mechanisms—can have a significant impact on long-term behavioral stability. This insight holds particular relevance for fields such as behavioral finance, public policy, and organizational decision-making, where managing time-inconsistent behaviors is essential. By formalizing the interplay between environmental noise, adaptive persistence, and subjective thresholds, the model not only captures existing empirical patterns but also provides a forward-looking framework for designing interventions aimed at fostering consistent long-term decision-making in uncertain environments.

The theoretical propositions articulated in the dynamic threshold model offer a comprehensive framework for understanding the intricate relationship between uncertainty, stochastic noise, and time preferences. While these propositions are firmly grounded in theoretical reasoning, their robustness and applicability are further reinforced by recent empirical studies. The following section examines key experimental and survey-based findings that align with and validate the model's core claims, demonstrating how real-world behavioral evidence supports the recursive mechanisms, adaptive persistence, and path-dependent dynamics central to the model.

6 Mechanisms in Action: Explaining Observed Phenomena

Recent empirical research provides strong support for the theoretical propositions embedded in the dynamic threshold model, validating its key insights into the role of stochastic shocks, adaptive persistence, and path-dependent behavior in shaping intertemporal preferences. These experimental and survey-based findings underscore the model's relevance in explaining observed behavioral patterns across diverse decision environments.

6.1 Proposition 1 (Noise Variance and Its Asymmetric Influence on Preference Switching)

The findings of Zhou, Li, and Liu (2021) provide crucial empirical support for Proposition 1 by examining the role of visual attention in shaping time preference reversals. Their research demonstrates that the allocation of attention across different attributes of a decision—such as reward magnitude, delay, and risk—can significantly influence whether individuals exhibit time-consistent or time-inconsistent preferences. This aligns with the proposition's emphasis on the relationship between perceived uncertainty and a dynamic subjective threshold. The study suggests that preferences are not rigid constructs but are instead highly susceptible to external environmental cues and the way information is processed. Attention, in this context, can be viewed as a mechanism influencing perceived uncertainty (z), thereby modulating the subjective dynamic threshold (θ_t). When attention is allocated differently, the level of perceived uncertainty shifts, altering the probability of preference reversals. This relationship underscores the recursive nature of the dynamic threshold model, where both internal and external stochastic factors interact to shape preference outcomes dynamically.

Importantly, Zhou, Li, and Liu (2021) demonstrate that attentional allocation introduces an inherent stochastic component into decision-making. Since attention is often reactive to external stimuli and subject to biases beyond deliberate control, its influence on decision-making cannot be entirely predicted. This aligns with the dynamic threshold model's proposition that stochastic noise amplifies uncertainty and exerts a disproportionate influence on preferences, especially near critical decision boundaries. The experimental findings illustrate how shifts in attentional focus alter perceived uncertainty, making preference switching more likely under certain conditions. This observation supports the core claim of Proposition 1, which posits that stochastic noise interacts with perceived uncertainty to drive preference reversals dynamically.

In a complementary study, Somasundaram and Eli (2022) demonstrate a significant interaction between risk and time preferences, showing that decision-makers' sensitivity to time delays is influenced by the probability of receiving a reward. These findings further reinforce the predictions of Proposition 1, highlighting that the subjective threshold (θ_t) is not static but evolves in response to external stochastic shocks, such as changes in risk levels. In the context of this research, variations in risk function as manipulations of perceived uncertainty (z), which in turn alter the dynamic threshold and influence time preferences. This direct interaction between risk

and time preferences aligns with the proposition's claim that switching probabilities depend on the relative position of uncertainty and the subjective threshold.

Similarly, Stanek et al. (2019) provide further empirical support by examining how uncertainty and reward processing interact to influence memory and choice behavior. Their results demonstrate that preferences are not merely products of intrinsic cognitive biases but are highly responsive to the external stochastic environment. The observed changes in choice and memory processes under varying levels of uncertainty can be interpreted through the dynamic threshold model as evidence of the recursive influence of stochastic noise on the subjective threshold. The findings suggest that preference reversals emerge not from fixed psychological parameters but from dynamic shifts in subjective uncertainty thresholds caused by external stochastic forces. This conclusion validates the core claim of Proposition 1, emphasizing that stochastic factors and environmental variability play a central role in driving preference instability and reversals.

Taken together, the studies by Zhou, Li, and Liu (2021), Somasundaram and Eli (2022), and Stanek et al. (2019) collectively support Proposition 1. They provide empirical evidence demonstrating that perceived uncertainty (z) and subjective thresholds (θ_t) are highly sensitive to stochastic factors, including shifts in attention, variations in risk, and changing reward structures. These findings underscore the recursive nature of the dynamic threshold model and its ability to capture preference dynamics in environments characterized by stochastic noise and fluctuating uncertainty.

6.2 Proposition 2 (Behavioral Instability Near the Threshold)

Proposition 2 builds on the insights of Proposition 1, emphasizing that behavioral instability intensifies near the subjective threshold and exhibits non-linear sensitivity to small changes in external stochastic factors. Zhou, Li, and Liu (2021) provide empirical validation for this proposition by demonstrating that visual attention patterns can significantly influence time preference reversals, particularly in situations where individuals are operating near critical decision boundaries. Their findings reveal that seemingly minor changes in attentional focus—whether on reward magnitude, time delay, or external stimuli—can result in substantial shifts in preferences. This aligns with the theoretical assertion that agents exhibit heightened sensitivity to stochastic noise when their decisions are made near the threshold. The dynamic threshold model predicts that even marginal fluctuations in perceived uncertainty or environmental noise can trigger behavioral reversals when agents are near this critical point.

The allocation of attention, therefore, may be seen as a mechanism that changes the subjective threshold (θ_t), thus leading to different behavioral responses. Their findings also demonstrate that choices are not always consistent over time and that small changes to external factors can have an out-sized impact on choice. This variability in choice may be interpreted as highlighting a key feature of the dynamic threshold, and that stochastic factors will influence choice, which is a key component of Proposition 2.

The study by Zhou, Li, and Liu (2021) also highlights how attention allocation acts as a stochas-

tic variable influencing preference dynamics. Since attention is often influenced by external environmental factors and is not entirely under the agent's conscious control, it introduces variability into decision outcomes. This observation supports the core claim of Proposition 2 by illustrating that preference instability is not merely a function of internal biases but is also amplified by external stochastic disturbances. The results suggest that decision boundaries become increasingly fragile as uncertainty intensifies, leading to greater behavioral volatility near the subjective threshold.

Furthermore, the findings of de Water et al. (2017) provide complementary insights by investigating the neural mechanisms underlying preference reversals. Their results reveal that time preferences are governed by non-linear interactions between different brain regions, particularly when decision-making occurs under uncertainty. This neural complexity aligns with the prediction of Proposition 2, which suggests that behavioral sensitivity near the threshold is inherently non-linear. The study emphasizes that agents operating close to their subjective threshold are more prone to external perturbations, resulting in time preference reversals and inconsistent choices.

The combined evidence from Zhou, Li, and Liu (2021) and de Water et al. (2017) supports the theoretical claims made in Proposition 2. Both studies demonstrate that preference instability is most pronounced near subjective thresholds, where stochastic noise amplifies behavioral sensitivity non-linearly. These findings reinforce the proposition's central claim that external stochastic forces and perceived uncertainty interact recursively near the decision threshold, creating heightened behavioral volatility.

In synthesizing these empirical observations, Proposition 2 emerges as a robust theoretical framework for understanding behavioral instability in dynamic decision environments. The heightened behavioral sensitivity observed near decision thresholds, as evidenced by the studies mentioned, validates the recursive and stochastic mechanisms articulated in the dynamic threshold model. These findings highlight the importance of considering both internal cognitive processes and external stochastic factors in explaining time preference reversals and behavioral instability.

6.3 Proposition 3 (Preference Instability Under Variable Noise)

Proposition 3 of the dynamic threshold model posits that under increasing levels of stochastic noise, choices become increasingly random, undermining long-term planning and stability. This proposition asserts that external stochasticity introduces variability into decision-making, disrupting consistent preference trajectories and amplifying randomness in behavioral outcomes. The findings of Somasundaram and Eli (2022) align with this claim, demonstrating that time and risk preferences interact dynamically, with sensitivity to time delays varying based on the probability of receiving a reward. Their results indicate that external stochastic factors, such as probabilistic risk, introduce additional layers of uncertainty (z), destabilizing preferences and making choices less predictable. This observation aligns with the core prediction of the dynamic threshold model, which emphasizes that increased stochasticity drives preference volatility by interacting recursively with subjective uncertainty thresholds.

The findings of Somasundaram and Eli (2022) further highlight that preferences are not static constructs but evolve dynamically in response to external stochastic stimuli. Variations in perceived risk amplify uncertainty, effectively increasing noise variance (σ^2), which destabilizes the subjective threshold (θ_t). This mechanism reinforces Proposition 3, suggesting that elevated levels of uncertainty disrupt agents' ability to maintain stable preferences over time. The results validate the claim that external stochasticity fundamentally alters the decision landscape, amplifying randomness in preference formation and increasing susceptibility to behavioral inconsistencies.

Complementing these observations, Stanek et al. (2019) explore the interaction between reward processing, memory, and environmental uncertainty. Their findings demonstrate that preferences are not solely determined by internal cognitive biases but are highly sensitive to external stochastic forces. Uncertainty distorts how individuals process reward-related information, introducing noise into preference structures and undermining stability. Within the framework of the dynamic threshold model, these distortions can be interpreted as increases in the variance of stochastic noise (σ^2), which destabilize the subjective threshold and lead to unpredictable decision outcomes. These results illustrate that preferences do not follow a fixed functional form but instead emerge from dynamic interactions between internal cognitive mechanisms and external stochastic disturbances.

Moreover, Stanek et al. (2019) emphasize that increased environmental noise disrupts the expected value of rewards, leading to behavioral randomness and inconsistent preferences. These findings align with the theoretical prediction of Proposition 3, which posits that higher stochastic noise levels will lead to greater randomness in choice behavior. The recursive structure of the dynamic threshold model captures this phenomenon by showing how accumulated uncertainty disrupts the stability of the subjective threshold, preventing agents from forming consistent long-term preferences.

Taken together, the studies by Somasundaram and Eli (2022) and Stanek et al. (2019) provide robust empirical support for Proposition 3. Both studies highlight that preferences are not static but are highly sensitive to external stochastic factors, with noise variance (σ^2) and perceived uncertainty (z) acting as key drivers of instability. The findings reinforce the recursive feedback mechanism central to the dynamic threshold model, demonstrating how elevated environmental stochasticity amplifies randomness, disrupts long-term preference trajectories, and increases the likelihood of unpredictable decision-making outcomes. These insights underscore the explanatory power of Proposition 3 in capturing the dynamic and stochastic nature of preference instability in uncertain decision environments.

6.4 Proposition 4 (Persistence, Stochasticity, and the Fragility of Long-Term Commitments)

Proposition 4 of the dynamic threshold model introduces the concept of path dependence, illustrating how accumulated external shocks and stochastic disturbances destabilize preferences over time, rendering long-term commitments fragile. The findings of Boutin, Petifour, and Megzari

(2023) provide compelling empirical support by analyzing the persistent effects of the COVID-19 pandemic on time preferences. Within the model, the pandemic serves as a series of historical uncertainty shocks (η_t), recursively influencing the subjective threshold and preventing preferences from stabilizing. The persistence of altered preferences post-pandemic suggests a high value for the persistence parameter (α), reinforcing the recursive feedback loop central to Proposition 4. These findings align with the theoretical prediction that cumulative stochastic shocks reshape preference trajectories and create long-term behavioral fragility.

Holden and Tilahun (2021) offer further validation by demonstrating how economic hardship and environmental shocks produce enduring changes in risk preferences. These shocks align with the model's historical uncertainty parameter (η_t), showing how repeated disturbances disrupt preference stability. The persistence parameter (α) explains why these changes endure over time, preventing preferences from reverting to pre-shock baselines. This evidence supports the recursive structure of the dynamic threshold model, emphasizing the cumulative impact of external stochastic forces on long-term decision dynamics.

Kuroishi and Sawada (2019) provide additional insights by showing how catastrophic disasters induce permanent changes in time preferences. Such events can be interpreted as large-scale historical uncertainty shocks (η_t), where their enduring impact corresponds to a high persistence parameter (α). The recursive adjustment mechanism of the dynamic threshold model explains why preferences remain disrupted long after the immediate shock subsides, reinforcing the fragility of long-term commitments under sustained uncertainty.

Giuliano and Spilimbergo (2024) extend these findings by analyzing the prolonged effects of aggregate economic shocks on individual preferences and belief systems. Their observations highlight how historical uncertainty shocks (η_t) interact with the persistence parameter (α), preventing a return to pre-shock stability. The recursive nature of the dynamic threshold model captures these persistent disruptions, showing how cumulative stochastic disturbances embed path dependence into preference trajectories.

Collectively, these studies provide robust empirical support for Proposition 4. The observed persistence of preference instability across diverse contexts validates the dynamic threshold model's prediction that historical uncertainty shocks (η_t) and the persistence parameter (α) interact recursively to destabilize long-term preferences. By integrating these mechanisms, the model offers a coherent explanation for why long-term commitments become fragile under prolonged uncertainty and why preferences remain sensitive to historical stochastic disturbances. This theoretical structure underscores the importance of accounting for path dependence and recursive feedback in understanding intertemporal decision dynamics.

6.5 Proposition 5 (Adaptive Persistence and Subjective Threshold Stability)

Proposition 5 of the dynamic threshold model introduces the concept of adaptive persistence, emphasizing that agents can strategically adjust the persistence parameter (α_t) to stabilize their subjective threshold (θ_t) in response to environmental volatility. This mechanism suggests a de-

gree of agency in managing preference stability, allowing agents to recalibrate their sensitivity to external stochastic disturbances over time.

The findings of Stanek et al. (2019) provide empirical support for this proposition by demonstrating that memory formation and reward processing are highly sensitive to environmental uncertainty. Their results highlight that neural mechanisms governing memory and reward evaluation are not static but adapt dynamically to environmental feedback. From the perspective of the dynamic threshold model, these adaptive neural changes can be interpreted as strategic adjustments to the persistence parameter (α_t), where agents effectively recalibrate how past experiences influence current decision-making processes. This interpretation aligns with the model's core assertion that agents can exercise control over persistence to optimize their ability to maintain threshold stability under varying levels of stochastic noise. Furthermore, Stanek et al. (2019) illustrate how improved consistency in memory mechanisms can contribute to a more stable subjective threshold, reinforcing the theoretical claim that preference stability is not merely an outcome of intrinsic cognitive biases but a result of adaptive calibration.

Complementing this, the research by Somasundaram and Eli (2022) demonstrates a significant interaction between time and risk preferences, revealing that agents adjust their sensitivity to time delays based on the probability of receiving a reward. This interaction implies that agents possess the capacity to balance competing priorities—risk sensitivity and temporal preferences—by dynamically managing persistence levels. Within the framework of the dynamic threshold model, these observed behavioral adjustments align with the role of the adaptive persistence parameter (α_t), which modulates the recursive adjustment of the subjective threshold in response to external stochastic signals. The observed flexibility in agent behavior supports the proposition's claim that adaptive persistence enables agents to recalibrate their decision-making processes strategically, enhancing stability in volatile environments while maintaining responsiveness to changing conditions.

Overall, these empirical findings reinforce the theoretical predictions of Proposition 5. Both studies illustrate that preference stability is not a fixed trait but a dynamic outcome shaped by adaptive adjustments to persistence (α_t) and subjective thresholds (θ_t). The dynamic threshold model provides a structural explanation for these observations, demonstrating how agents navigate external stochastic environments by recalibrating persistence levels to balance stability and adaptability. This theoretical framework emphasizes the recursive nature of preference formation and highlights the active role of agents in maintaining long-term behavioral stability amid environmental uncertainty.

6.6 Proposition 6 (Persistence and the Speed of Subjective Threshold Recovery)

Proposition 6 of the dynamic threshold model highlights a fundamental trade-off between the stability of the subjective threshold (θ_t) and the speed of recovery following stochastic shocks. This proposition emphasizes that agents may prioritize either stability or responsiveness, depending on their preference structure and environmental conditions. The experimental findings of Zhou,

Li, and Liu (2021) provide empirical validation for this trade-off by demonstrating how variations in visual attention influence time preference reversals. Their results suggest that patterns of attentional focus can be mapped onto different levels of the persistence parameter (α), with higher persistence corresponding to a more stable focus of attention and lower persistence associated with frequent shifts in focus. This empirical observation aligns with the theoretical prediction that higher persistence enhances stability but reduces sensitivity to external changes, while lower persistence allows for faster adaptation but amplifies volatility in preferences. These dynamics illustrate how behavioral adjustments after stochastic shocks are shaped by the interplay between persistence and environmental volatility, reinforcing the core argument of Proposition 6.

Additionally, the findings from Zhou, Li, and Liu (2021) suggest that time preference reversals can serve as empirical indicators of the speed of subjective threshold recovery. The observed behavioral fluctuations following changes in attentional patterns reflect the recursive adjustment mechanism central to the dynamic threshold model. When agents exhibit rapid reversals in time preferences, it implies a lower persistence parameter (α), enabling quicker recovery at the cost of increased instability. Conversely, more stable preferences, accompanied by slower recovery, indicate a higher persistence level. These findings underscore the model's assertion that the trade-off between stability and recovery speed is a defining characteristic of dynamic decision environments influenced by stochastic shocks.

The review by Lempert and Phelps (2016) further supports the theoretical underpinnings of Proposition 6 by emphasizing the malleability of intertemporal choices. Their work highlights that preferences are not fixed but are continually reshaped by both internal cognitive mechanisms and external environmental factors. This aligns with the model's argument that persistence (α) and threshold stability (θ_t) are dynamic parameters that evolve in response to changing external conditions. Lempert and Phelps also illustrate how agents may strategically adapt their decision-making parameters over time, providing empirical context for understanding the flexibility embedded within the dynamic threshold model. The capacity for agents to recalibrate their persistence parameter in response to external volatility aligns with the model's prediction that preference stability and adaptability are not mutually exclusive but exist along a continuum defined by environmental demands and strategic adjustments.

Collectively, these studies reinforce the theoretical claims of Proposition 6. They demonstrate how persistence acts as a regulatory mechanism, balancing the need for stability with the capacity for rapid recovery. The empirical findings from Zhou, Li, and Liu (2021) and the theoretical insights from Lempert and Phelps (2016) validate the recursive interaction between persistence and environmental volatility, illustrating how agents dynamically navigate preference stability and responsiveness under stochastic conditions. This trade-off, central to the dynamic threshold model, provides a robust explanation for the observed variability in behavioral outcomes across different decision contexts.

The theoretical propositions and their empirical validations presented in the preceding sections collectively establish the dynamic threshold model as a robust framework for understand-

ing time preferences under stochastic uncertainty. By capturing the recursive interactions between subjective thresholds, stochastic shocks, and adaptive persistence, the model provides a nuanced explanation of both short-term behavioral fluctuations and long-term preference fragility. With these insights in place, the following section situates the model within the broader literature, discusses its contributions and limitations, and outlines potential directions for future research.

7 Discussion and Limitations

This study presents a dynamic threshold and uncertainty model, offering a novel theoretical framework for understanding time preferences. The core mechanism, a dynamic subjective threshold influenced by perceived uncertainty and external shocks, challenges the assumption of fixed preferences, as in models with fixed present bias and discount factors, or those that assume a static discounting function. This work highlights that preferences are not simply a function of internal bias, but are also determined by stochasticity in the environment, and a response to external shocks, which all have a cumulative impact over time. The analyses has also demonstrated how this framework can be used to derive novel theoretical results about the relationship between noise, thresholds and preference change, as well as how individuals can adapt their persistence parameter to create more stable preferences.

The model challenges assumptions of fixed preferences, and instead proposes that choice is highly sensitive to the external environment, that this sensitivity will increase when external uncertainty approaches a dynamic threshold, and that these fluctuations may have a long term impact by changing the long term value of the threshold. By demonstrating that internal beliefs and external stochasticity can destabilize preferences over time, and that past history plays a role in the formation of preferences, the model goes beyond standard discounting functions and emphasizes dynamic mechanisms for choice. The model also demonstrates that choices are not simply a function of a fixed present bias, but are influenced by dynamic interactions between a threshold, past shocks, and external stochasticity.

The framework of Halevy (2015), which defines the normative features of a time consistent choice, is useful for highlighting the nature of the deviations from consistent choice that are captured by this work. The literature reviewed by Ericson and Laibson (2019) highlights the diversity of the approaches that are used, and also emphasizes the importance of new theoretical perspectives. Further, the results in the works of Harris and Laibson (2001), as well as Gabaix (2019), demonstrate that there may be important limitations in some of the existing approaches that are used to model decision making. The key limitation of these, when compared with the work presented, is that they often assume fixed parameters for choice, rather than highlighting that these parameters, and preferences, may change dynamically over time. This is a core idea that is missing from existing literature.

However, the model does have limitations. The model has been deliberately constrained by certain simplifying assumptions, which, while useful for theoretical tractability, may not ade-

quately represent the full complexity of human decision-making. One key limitation is the simplification of all external influences as either stochastic shocks, or as a source of environmental uncertainty. Although this is useful for highlighting the impact that external sources of information may have, it does not capture how other processes, such as cognitive limitations or inattentional biases, may influence choices. Additionally, although some of the key components of the model have been supported by experimental and empirical data, the model has not, as yet, been empirically tested using an experimental design that is built to specifically address the core mechanisms that it has introduced. This lack of direct empirical evidence represents a clear limitation and an area for future research. Since the framework has been proposed to enhance the understanding of real-world behavior, validation through empirical data remains essential.

Future research should seek to extend this model by explicitly considering the potential for cognitive limitations, attentional biases, and other real-world constraints on human choices. The role of memory may also be an important mechanism that can be included. Further studies should also explore how to test the model empirically, specifically targeting the hypothesized interaction between dynamic thresholds and uncertainty, and exploring how stochasticity and past shocks may influence both current and future decisions. This can be done by developing experimental paradigms that explicitly manipulate the stochastic nature of the decision-making environment, and by testing how preferences change as a result of such interventions, allowing for a more explicit test of the model's claims. Such research can also extend the approach to include strategic interactions in games, as well as explore how the adaptive mechanisms of the model may be used to better understand choices.

The dynamic threshold and uncertainty model offers a new way of thinking about intertemporal choice by incorporating dynamic and stochastic forces as core mechanisms of choice. This approach builds upon and extends previous work and has the potential to offer novel insights that go beyond the constraints that are placed on traditional models, and it is hoped that this will be a useful stepping stone for future researchers in this field.

8 Conclusion

The dynamic threshold model presents a comprehensive theoretical framework for understanding time preferences by incorporating stochastic shocks, adaptive persistence, and recursive subjective thresholds. Moving beyond traditional static models like the β-δ discounting framework, this model emphasizes the dynamic and context-sensitive nature of decision-making under uncertainty. Its recursive structure captures behavioral instability near subjective thresholds, long-term path dependence caused by accumulated stochastic shocks, and the intricate trade-off between stability and responsiveness in adaptive decision environments.

In the short term, the model elucidates how heightened uncertainty and stochastic noise amplify preference reversals near decision thresholds, providing a structural explanation for observed behavioral volatility. Over the long term, it reveals how cumulative stochastic shocks

embed persistent instability into preference trajectories, creating path dependence that prevents reversion to stable baselines. Additionally, the adaptive persistence parameter (α) emerges as a pivotal mechanism, balancing resilience to volatility with responsiveness to external changes.

By addressing both short-term fluctuations and long-term fragility, the dynamic threshold model bridges limitations in traditional frameworks. Its theoretical propositions align closely with empirical findings, offering clarity on phenomena such as cognitive processing delays, behavioral instability near uncertainty thresholds, and the enduring effects of environmental shocks.

Beyond theoretical contributions, the model offers actionable insights for experimental designs, behavioral interventions, and policy frameworks aimed at stabilizing preferences in uncertain decision contexts. As recent work by Laibson and Maxted (2023) has highlighted the need for methodological rigor when simulating present-biased behaviors, it is also important to understand the underlying mechanisms that may give rise to time-inconsistent choices. By formalizing the effect of stochastic shocks and noise on the subjective threshold, this model offers a complementary framework for understanding the potential causes of preference instability that have been observed in simulation. Future research should focus on refining empirical tests of the model's predictions, exploring its applications in strategic environments, and extending its utility to related domains of intertemporal choice.

In conclusion, the dynamic threshold model provides a flexible and robust foundation for advancing the understanding of time preferences, offering both theoretical rigor and practical relevance in addressing the complexities of decision-making under uncertainty.

Declarations

Funding: Not applicable.
Conflict of interest: The author declares no conflict of interest.
Ethical approval: Not applicable.

References

[1] Akerlof, G. A. (1991). Procrastination and obedience. *The American Economic Review, 81*(2), 1–19.

[2] Alós-Ferrer, C., Jaudas, A., & Ritschel, A. (2021). Preference reversals and attention in intertemporal choice: Evidence from eye-tracking. *Journal of Behavioral Decision Making, 34*(2), 196–212. https://doi.org/10.1002/bdm.2208

[3] Andersson, O., Holm, H. J., Tyran, J.-R., & Wengström, E. (2016). Risk aversion relates to cognitive ability: Preferences or noise? *Journal of the European Economic Association, 14*(5), 1129–1154. https://doi.org/10.1111/jeea.12179

[4] Ashby, N. J. S., Dickert, S., & Glöckner, A. (2012). Attention and preference in risky choice: The influence of visual fixation patterns on decision-making. *Judgment and Decision Making, 7*(3), 267–276. https://doi.org/10.1016/j.jecp.2011.11.012

[5] Ashby, N. J. S., Walasek, L., & Glöckner, A. (2015). Response dynamics in intertemporal choice. *Cognition, 136*, 43–60. https://doi.org/10.1016/j.cognition.2014.11.002

[6] Blavatskyy, P. R. (2022). Intertemporal choice as a tradeoff between cumulative payoff and average delay. *Journal of Risk and Uncertainty, 64*(1), 89–107. https://doi.org/10.1007/s11166-022-09378-4

[7] Boutin, D., Petifour, L., & Megzari, H. (2023). Permanent instability of preferences after COVID-19 crisis: A natural experiment from urban Burkina Faso.

[8] Carvalho, L. S., Meier, S., & Wang, S. W. (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review, 106*(2), 260–284. https://doi.org/10.1257/aer.20140481

[9] Cohen, J., Ericson, K. M., Laibson, D., & White, J. M. (2020). Measuring time preferences. *Journal of Economic Literature, 58*(2), 299–347. https://doi.org/10.1257/jel.20191558

[10] DellaVigna, S. (2018). Structural behavioral economics. In *Handbook of Behavioral Economics: Applications and Foundations 1* (Vol. 1, pp. 613–723). North-Holland.

[11] de Water, E., Mies, G. W., Figner, B., Yoncheva, Y., van den Bos, W., Castellanos, F. X., ... & Scheres, A. (2017). Neural mechanisms of individual differences in temporal discounting of monetary and primary rewards in adolescents. *NeuroImage, 153*, 198–210. https://doi.org/10.1016/j.neuroimage.2017.03.051

[12] Ericson, K. M., & Laibson, D. (2019). Intertemporal choice. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 2* (pp. 1–67). North-Holland/Elsevier. https://doi.org/10.1016/bs.hesbe.2018.12.001

[13] Franco-Watkins, A. M., Mattson, R. E., & Jackson, M. D. (2016). Now or later? Attentional processing and intertemporal choice. Journal of Behavioral Decision Making, 29(2–3), 206–217. https://doi.org/10.1002/bdm.1913

[14] Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*(2), 351–401. https://doi.org/10.1257/002205102320161311

[15] Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1* (pp. 191–239). Elsevier.

[16] Giuliano, P., & Spilimbergo, A. (2024). Aggregate shocks and the formation of preferences and beliefs.

[17] Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica, 83*(1), 335–352. https://doi.org/10.3982/ECTA12304

[18] Harris, C., & Laibson, D. (2001). Dynamic choices of hyperbolic consumers. *Econometrica, 69*(4), 935–957. https://doi.org/10.1111/1468-0262.00226

[19] He, X., & Bhatia, S. (2018). Stochastic variability in intertemporal choice. *Journal of Economic Behavior & Organization, 148*, 180–193. https://doi.org/10.1016/j.jebo.2018.02.001

[20] Holden, S. T., & Tilahun, M. (2021). Shocks and stability of risk preferences (No. 05/21). Centre for Land Tenure Studies Working Paper.

[21] Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). Noise: A flaw in human judgment. Hachette UK.

[22] Kuroishi, Y., & Sawada, Y. (2024). On the stability of preferences: Experimental evidence from two disasters. *European Economic Review, 161*, 104632. https://doi.org/10.1016/j.euroecorev.2024.104632

[23] Laibson, D. I. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics, 112*(2), 443–478. https://doi.org/10.1162/003355397555253

[24] Laibson, D. I., & Maxted, P. (2023). The beta-delta-delta sweet spot. *NBER Working Paper No. w30822*. https://ssrn.com/abstract=4320570

[25] Lee, D. (2013). Decision making: From neuroscience to psychiatry. *Neuron, 78*(2), 233–248. https://doi.org/10.1016/j.neuron.2013.04.008

[26] Lempert, K. M., & Phelps, E. A. (2016). The malleability of intertemporal choice. *Trends in Cognitive Sciences, 20*(1), 64–74. https://doi.org/10.1016/j.tics.2015.09.005

[27] Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics, 107*(2), 573–597.

[28] Mitchell, S. H., Wilson, V. B., & Karalunas, S. L. (2015). Comparing hyperbolic, delay-amount sensitivity and present-bias models of delay discounting. *Behavioural Processes, 114*, 52–62. https://doi.org/10.1016/j.beproc.2015.02.003

[29] Oberrauch, L., & Kaiser, T. (2022). Cognitive ability, financial literacy, and narrow bracketing in time-preference elicitation. Journal of Behavioral and Experimental Economics, 98, 101844. https://doi.org/10.1016/j.socec.2022.101844

[30] O'donoghue, T., & Rabin, M. (2001). Choice and procrastination. The Quarterly Journal of Economics, 116(1), 121–160. https://doi.org/10.1162/003355301556365

[31] Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty, 23*(1), 5–32. https://doi.org/10.1023/A:1011198414681

[32] Samuelson, P. A. (1937). A note on measurement of utility. The Review of Economic Studies, 4(2), 155–161. https://doi.org/10.2307/2967612

[33] Schildberg-Hörisch, H. (2018). Are risk preferences stable? Journal of Economic Perspectives, 32(2), 135–154. https://doi.org/10.1257/jep.32.2.135

[34] Skylark, W. J., Farmer, G. D., & Bahemia, N. (2021). Inference and preference in intertemporal choice. Judgment and Decision Making, 16(2), 422–459.

[35] Somasundaram, J., & Eli, V. (2022). Risk and time preferences interaction: An experimental measurement. *Journal of Risk and Uncertainty, 65*, 215–238. https://doi.org/10.1007/s11166-022-09394-9

[36] Stanek, J. K., Dickerson, K. C., Chiew, K. S., Clement, N. J., & Adcock, R. A. (2019). Expected reward value and reward uncertainty have temporally dissociable effects on memory formation. *Journal of Cognitive Neuroscience, 31*(10), 1443–1454. https://doi.org/10.1162/jocn_a_01411

[37] Stewart, N., Hermens, F., & Matthews, W. J. (2016). Eye movements in risky choice. Journal of Behavioral Decision Making, 29(2–3), 116–136. https://doi.org/10.1002/bdm.1903

[38] Sui, X. Y., Liu, H. Z., & Rao, L. L. (2020). The timing of gaze-contingent decision prompts influences risky choice. Cognition, 195, 104077. https://doi.org/10.1016/j.cognition.2019.104077

[39] Zhou, Y. B., Li, Q., & Liu, H. Z. (2021). Visual attention and time preference reversals. *Judgment and Decision Making, 16*(4), 1010–1038.

Appendix A: Derivation of the Equilibrium Value for the Subjective Threshold

This appendix presents a detailed derivation of the equilibrium value for the subjective threshold (θ_t). This equilibrium, or long-run mean, serves as a baseline against which to compare the dynamic fluctuations caused by noise, shocks, and the persistence parameter discussed in the main text.

To derive a stable equilibrium value, it becomes necessary to posit some assumptions regarding the behavior of the variables within the model over extended time periods. It is assumed that the expected value of perceived uncertainty (z_t) converges to a constant value over time, indicating that the agent's overall level of uncertainty concerning the outcome of long-term and short-term rewards tends towards a stable mean. Mathematically, this may be expressed as $\lim_{t\to\infty} \mathbb{E}[z_t] = z^*$, where z^* denotes a constant. Furthermore, it is assumed that the expected value of the exogenous shocks (η_t) is zero over the long run, such that $\mathbb{E}[\eta_t] = 0$. This assumption facilitates the isolation of the effect of perceived uncertainty upon the equilibrium. A final assumption is that, in the absence of external fluctuations, the subjective threshold reaches a steady state where its expected value remains constant over time. This condition is expressed as $\lim_{t\to\infty} \mathbb{E}[\theta_{t+1}] = \lim_{t\to\infty} \mathbb{E}[\theta_t] = \theta^*$, where θ^* represents the equilibrium value of the subjective threshold.

Recall that the subjective threshold evolves according to the following recursive equation (Equation 1 from the main text):

$$\theta_{t+1} = \alpha\theta_t + (1-\alpha)z_t + (1-\alpha)\epsilon_t + \eta_t$$

Where θ_t denotes the subjective threshold at time t, α is the persistence parameter, z_t represents perceived uncertainty at time t, ϵ_t is a noise term such that $\epsilon_t \sim N(0, \sigma^2)$, and η_t is an exogenous shock at time t. Taking the expectation of both sides yields:

$$\mathbb{E}[\theta_{t+1}] = \mathbb{E}[\alpha\theta_t + (1-\alpha)z_t + (1-\alpha)\epsilon_t + \eta_t]$$

Utilizing the linearity of the expectation operator, this equation is reformulated as:

$$\mathbb{E}[\theta_{t+1}] = \alpha\mathbb{E}[\theta_t] + (1-\alpha)\mathbb{E}[z_t] + (1-\alpha)\mathbb{E}[\epsilon_t] + \mathbb{E}[\eta_t]$$

Given that the expected value of the noise and shock terms are zero, such that $\mathbb{E}[\epsilon_t] = 0$ and $\mathbb{E}[\eta_t] = 0$, the expression can be simplified to:

$$\mathbb{E}[\theta_{t+1}] = \alpha\mathbb{E}[\theta_t] + (1-\alpha)\mathbb{E}[z_t]$$

Based on the assumptions previously outlined, the equilibrium value of the threshold may be derived. In the long run, as t approaches infinity, the expected values of θ_t and z_t will converge to

θ^* and z^*, respectively. Consequently, the recursive equation may be rewritten as:

$$\theta^* = \alpha\theta^* + (1 - \alpha)z^*$$

Here, the long-run expected values for the threshold (θ^*) and the perceived uncertainty (z^*) have been substituted. Algebraic manipulation can then be used to solve for θ^*:

$$\theta^* - \alpha\theta^* = (1 - \alpha)z^*$$

$$\theta^*(1 - \alpha) = (1 - \alpha)z^*$$

$$\theta^* = \frac{(1 - \alpha)z^*}{(1 - \alpha)}$$

This reduces to:

$$\theta^* = z^*$$

The derived result demonstrates that the equilibrium value of the subjective threshold (θ^*) is equal to the long-run expected value of perceived uncertainty (z^*). This equilibrium value functions as a crucial baseline against which dynamic changes in the threshold may be evaluated. If an agent's threshold value at time t were significantly high, relative to the equilibrium threshold, it might be expected that the threshold will decline over time, subject to stochastic shocks. Conversely, should the current value of the threshold be less than the equilibrium, the long-term trend would be toward an increase in value, again subject to stochastic variance. The equilibrium represents the point toward which the subjective threshold will tend in the absence of shocks or random noise, highlighting a key stability point within the model. Additionally, the equality of the equilibrium value of the threshold and the expected long-run uncertainty demonstrates that an agent's perception of uncertainty is central to establishing the benchmark for intertemporal choice. Increased long-run uncertainty associated with an event or behavior will lead to an increase in the subjective threshold, rendering long-term investments less appealing.

This derivation does not imply that agents always achieve the equilibrium threshold; rather, it provides a baseline toward which the threshold will gravitate in the absence of stochastic noise. Deviations from the equilibrium are central to this model and are a consequence of the combined effects of the persistence parameter, random noise, and exogenous shocks. The existence of the equilibrium underscores the importance of stochastic noise and shocks. Any deviation of the threshold from this long-run mean value may be attributed to the stochastic elements within the system. Although the equilibrium threshold is not dependent on the persistence parameter, this parameter is central to understanding the dynamics of the system. Higher persistence implies slower convergence toward the equilibrium threshold and reduced responsiveness to new information, while a lower persistence parameter will produce a faster response to shocks. The adaptive model, characterized by lower persistence, is therefore expected to approach the equilibrium threshold more rapidly.

It should be noted that the validity of the derived equilibrium is contingent upon the stated assumptions. The assumption that perceived uncertainty tends towards a constant value may not always hold, meaning there may not always be an equilibrium. It could be useful to consider cases where perceived uncertainty is itself a stochastic variable, or where the distribution of values shifts over time. Moreover, this equilibrium does not imply that individuals are maximizing any objective function, rather, it should be considered as a baseline for comparisons. Future investigations might consider whether this threshold could be treated as an optimizing process with respect to a specified objective function. This study is focused on long-run equilibrium, and future analysis could be devoted to exploring different conditions in the short run, and what properties are required for a system to reach equilibrium values.

In conclusion, the derivation presented in this appendix provides a concrete understanding of the long-run behavior of the subjective threshold. The result that the subjective threshold will tend toward the long-run expected value of perceived uncertainty, allows comparisons between the actual value of the threshold and its equilibrium, which further strengthens the interpretation of the main results. Furthermore, while the threshold will tend toward its equilibrium value over time, it is likely to never fully reach it due to the constant influence of stochastic elements and the degree of inertia associated with the persistence parameter. The derived equilibrium provides a baseline value for this study and it can also serve as a valuable reference for future research.

Appendix B: Derivation of the Adaptive Persistence Parameter

The persistence parameter (α_t) represents the extent to which past subjective thresholds (θ_t) influence future thresholds. In many recursive models, this parameter is considered constant over time. However, when agents adapt their persistence based on observed deviations from an optimal reference point (θ_0), it becomes necessary to express α_t as a dynamic function of these deviations.

To ensure stability in the presence of large threshold deviations, the functional form is expressed as:

$$\alpha_t = \alpha_0 - \gamma(\theta_t - \theta_0)^2$$

The parameter $\gamma > 0$ represents the sensitivity of persistence to threshold deviations. Larger values of γ imply that even small deviations from the optimal threshold (θ_0) lead to significant reductions in persistence (α_t). Conversely, smaller values of γ indicate that persistence remains relatively stable, even when deviations from the optimal threshold are moderate.

This functional form ensures that persistence decreases in environments where the subjective threshold experiences substantial volatility, thereby preventing excessive reliance on historical benchmarks under unstable conditions. It also facilitates a self-correcting mechanism where stability is promoted when deviations decrease.

The need for adaptive persistence arises because large deviations from the optimal threshold suggest misalignment with stable benchmarks. In such situations, maintaining a high persistence parameter (α_t) risks entrenching agents in suboptimal trajectories. Conversely, when the thresh-

old remains close to its optimal level, high persistence ensures stability and smooth adjustments over time. Therefore, the persistence parameter must decrease when deviations from the optimal reference point become significant, allowing the system to respond more effectively to uncertainty and shocks.

To construct an appropriate functional form for adaptive persistence, the following qualitative behavior must hold. Persistence (α_t) should decrease as the deviation ($\theta_t - \theta_0$) increases. This adjustment should be symmetric for positive and negative deviations. Small deviations should lead to minimal changes in persistence, preserving stability. Finally, the relationship should exhibit smooth and continuous behavior.

A quadratic penalty function naturally satisfies these requirements due to its symmetry and ability to penalize larger deviations more heavily. The adaptive persistence parameter can therefore be expressed as:

$$\alpha_t = \alpha_0 - \gamma \left(\theta_t - \theta_0\right)^2$$

In this expression, $\alpha_0 > 0$ represents the baseline level of persistence when the subjective threshold remains close to the optimal reference point. The parameter $\gamma > 0$ measures the sensitivity of the persistence parameter to deviations from the optimal threshold. The term $(\theta_t - \theta_0)^2$ captures the magnitude of these deviations symmetrically and ensures that large deviations reduce persistence more significantly.

The quadratic form emerges from a penalty framework where deviations from optimal behavior incur increasing costs at an accelerating rate. Formally, this can be framed within a loss function associated with persistence:

$$L(\alpha_t) = \alpha_0 - f(\theta_t, \theta_0)$$

To ensure symmetry and smooth adjustments, the penalty function $f(\theta_t, \theta_0)$ depends quadratically on deviations:

$$f(\theta_t, \theta_0) = \gamma \left(\theta_t - \theta_0\right)^2$$

Substituting this penalty function into the loss framework yields:

$$\alpha_t = \alpha_0 - \gamma \left(\theta_t - \theta_0\right)^2$$

This functional form captures the desired adaptive behavior. When deviations from the optimal threshold are small, the penalty term becomes negligible, and persistence remains close to its baseline value (α_0). As deviations grow larger, persistence decreases systematically, reducing reliance on historical thresholds and allowing the system to recalibrate more effectively.

The derived expression balances stability and responsiveness in environments characterized by stochastic variability and misalignment from optimal benchmarks. It provides a mathematically tractable and behaviorally interpretable model for adaptive persistence in dynamic decision-making processes.