A Neuro-Autopilot Theory of Habit: Evidence from Canned Tuna

Ryan Webb
Univ. of Toronto

Clarice Zhao
McGill

Matthew Osborne
Univ. of Toronto

Peter Landry
Univ. of Toronto

Colin Camerer
Caltech *

Feb 2022

Abstract

We integrate a neuroeconomic concept of habit into a consumer choice model. We propose that habit represents a distinct decision-making mode in which past choices are automatically repeated, in contrast with state-dependent utility maximization; transitions between these decision modes are governed by the reliability of a reinforcement learning algorithm, such that habits arise when the choice environment is sufficiently stable. We estimate and test this model on product choice in the canned tuna category between 2006 and 2010, a period of considerable price and product variation which included a package down-sizing event. We find that a substantial proportion of choice persistence is due to a habitual automation of consumption, in addition to a degree of state-dependent utility.

*We would like to thank A. Ching, G. Compiani, J.P. Dube, M. Pycia, O. Urminsky, S. Zhang, W. Wood and seminar participants at Chicago Booth, the University of Zurich, the 2021 Marketing Science Conference, the 2022 Lake Arrowhead Decision Neuroscience meeting and the 2022 SOUR BEER conference for helpful comments. Funding was provided by the Sloan Foundation G-2018-1125 (CFC) and a SSHRC Insight Development Grant.

Disclaimers:

Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business.

The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Intro

Habits are an important feature of human behaviour. They are not merely prevalent. There is also a wide and implicit recognition of their significance to our well-being, for better or worse. As noted by Gary Becker, "...the main reason habitual behavior permeates most aspects of life is that habits have an advantage in the biological evolution of human traits" (Becker, 1996, pg. 9).

A hallmark of a habit is that behaviour is persistent over time. In economic settings, this manifests as the empirical observation that present consumption often increases with past consumption. The predominant theoretical explanation for this persistence is inter-temporal complementarity in preferences, whereby a consumer's (marginal) utility for a good increases with past consumption of the good (e.g. Pollak, 1970; Ryder and Heal, 1973; Becker and Murphy, 1988; Crawford, 2010). This explanation also underlies an empirical literature examining 'state-dependent utility' in consumer choice — modeled as a positive effect of past consumption on current utility (Heckman, 1981; Keane, 1997; Shum, 2004; Seetharaman, 2004; Dubé et al., 2010; Thomadsen and Seetharaman, 2018; Kong et al., 2022).[1] Such persistence has been documented over a wide variety of product categories with important welfare implications because it provides a means for market leaders to leverage market-share. For example, Bronnenberg et al. (2012) find substantial persistence in the brands consumers buy across geographic regions, with 40% of variation in

[1]When paired with an assumption of rational expectations, state-dependent utility yields a consumer who foresees their habit formation and adjusts their consumption accordingly, i.e. a 'rational addiction' as proposed by Becker and Murphy (1988). Empirical tests of this claim are reported by Laporte et al. (2017); Auld and Grootendorst (2004); Olekalns and Bardsley (1996). As discussed by Landry (2019), however, the "adjacent complementarity" on which rational addiction theory is based is not empirically supported — only "distant complementarity" over longer time horizons is observed — in the realm of addictive consumption.

market shares due to persistent brand preferences and only 60% due to new supply conditions.[2]

While a state-dependent utility model can capture choice persistence, we argue that it is not a sufficient definition of habit. As one example, a state-dependent utility model explicitly assumes that prices and other attributes in the utility specification are collected and compared on every choice occasion. This assumption nullifies one of the primary benefits of a habit: that seemingly complex actions can be automated with minimal cognitive resources. It implies that all consumption occasions, even habitual ones, are utility-maximizing, and that a single measure of price elasticity is sufficient for capturing demand changes no matter how small or large the price change.

We study a model of habitual consumption behavior based on findings in psychology and neuroscience, and compare it to the existing state-dependent utility approach. Despite decades of research, there has been minimal integration of psychological habit research into economic discourse. In psychology, habits are considered a specific form of automatic behaviour which is directly cued by contexts that have been learned to be rewarding over time (Siegel et al., 1982; Wood and Neal, 2009). Building on this definition, Landry et al. (2021) introduce a theory of habitual choice based on the finding in neuroscience that the reliability of 'reward prediction errors' guide human and animal learning.[3]

[2]The degree of persistence is weaker for people who moved when they were younger. Consistent with this finding, brand loyalty is generally shown to be stronger among older consumers (Lambert-Pandraud & Laurent, 2010). Bronnenberg et al. (2012) report an estimate of the utility weight of past consumption vs. current attributes of 37%.

[3]A foundational empirical result in neuroeconomics is the existence of a reward prediction error (RPE) signal encoded via dopamine neurons (Schultz et al., 1997). This signal compares rewards to expectations and forms the basis of a reinforcement learning computation in instrumental and conditioned learning tasks. Lee et al. (2014) demonstrate the existence of an unsigned RPE signal which arbitrates between conditioned responses and more goal-directed planning systems.

Two modes of decision-making are proposed: a "habitual" mode in which the previous choice is automatically repeated, and a "model-based" mode in which utility is maximized using all available information.[4] To arbitrate between these systems, the consumer forms utility predictions and tracks their reliability. The consumer enters a habit when utility predictions are sufficiently reliable (i.e. when choice outcomes match predictions) and exits habit mode when there is sufficient doubt about these utility predictions.

One key prediction of Landry et al. (2021) is that small shocks to utility, or equivalently small price changes, can be insufficient to jolt a consumer out of their habit. When the choice environment is relatively stable, habits will be formed quickly. However large utility shocks (or relatively unstable prices) will cause consumers to re-optimize their consumption patterns. The theory can therefore predict that the price elasticity of demand differs depending on the size of the price change. It also implies periods of seemingly "sub-optimal" persistent behaviour that is rectified after large shocks to the choice environment.

In this paper, we adapt the Landry et al. (2021) theory of habit formation as an empirical consumer choice model and apply it to analyze a tumultuous period for the canned tuna industry. In 2008, canned tuna producers shrunk the size of their cans in the US market from 6oz to 5oz, with the introduction of the 5oz cans staggered across brands and stores. This downsizing event occurred in a product category that, in previous years, had been relatively stable, and preceded a significant adjustment in prices and relative market share between

[4]The modelling assumption that habit-based choice is completely "automatic" while model-based choice is perfectly rational is intended to highlight the distinction between the two systems as conceptualized in the psychology literature (Wood et al., 2021). In practise, the model-based system might not appear "perfectly rational" due to the presence of search costs or other informational issues which contribute to the agent mis-identifying the seemingly optimal choice (as in case of the Brazilian beer market or the London Tube examples, see footnote 6).

brands. Both factors make it ideal to test a model of how fluctuations and instability in the choice environment can lead to broken habits. We therefore estimate a structural model of habitual choice that provides a direct test for habitual decision-making as well as an estimate of the proportion of consumers who are habitual in any given period, and compare this model to the state-dependent utility account of choice persistence.

We find that a considerable degree of choice persistence in canned tuna consumption is due to habitual automation of choice behaviour. While we do still observe a degree of state-dependent utility, it is roughly 40% smaller in magnitude when we allow for habitual autopilot. This suggests that the welfare implications of persistence is more pronounced than previously thought, and provides a stronger justification for policies that increase variation in the choice environment. Overall, we estimate that 12% of consumers are in a habit before the can size change, dropping to 10% during the can introduction, with roughly 17% of habitual consumers exiting a habit during the new can introduction.

Our study relates to a number of different literatures in the social sciences. One recent literature addresses the identification of boundedly rational choice models that can yield choice persistence. Cerigioni (2021) studies the identification of dual-process models in which choices are automatically repeated when choice environments are sufficiently "similar" to each other, motivated from cognitive psychology research on perceptual fluency. He shows this similarity function is identified from choices, in principle, under the assumption that it is shared across a large population of consumers. By contrast, our doubt stock is defined over an accumulation of utility prediction errors over time, consistent with the observation that habits are learned via reward associations and are not formed immediately.[5] Our neuro-autopilot model is also related to limited consideration models in behavioural decision theory (Masatlioglu et al., 2012; Manzini and Mariotti, 2014), in particular the class of "default specific consideration" models studied by Abaluck and Adams-Prassl (2021). The difference is that the automated choice reflects past learned utilities, rather than a default. Our model can be similarly identified from own and cross-price elasticity asymmetries.

Matysková et al. (2020) show that persistent choices can be an optimal response to costly information acquisition, using the dynamic rational inattention framework of Steiner et al. (2017). In a highly-controlled lab experiment, subjects demonstrate persistence in their choices when states of the world are serially correlated, and this persistence disappears when states are independent. On a grander scale, Giuliano and Nunn (2021) argue that persistence in cultural norms are transmitted over generations in societies where the environment is more stable. Our data provide evidence for this trade-off at the level of an individual consumer: individual behaviour becomes persistent when the choice environment is stable and more sensitive to observables (like price) when it is more unstable.

Finally, a number of papers examine consumer preference in settings where a period of forced experimentation leads to an apparent improvement in consumer welfare. For example, Larcom et al. (2017) find that a temporary two-day labour strike forced some London tube riders to explore alternative commutes,

[5] Like Cerigioni (2021), we understand that both environmental context and reward processes are important for forming (and breaking) habits. The cognitive psychology and neuroscience literature has also long emphasized the importance of contextual cues (e.g. locations and social groups) in habit formation (Wood and Neal, 2009; Wood et al., 2021). Similarly, the number and frequency of learned reward associations is a key determinant of habitual behaviour in animals and humans (Adams and Dickinson, 1980; Pool et al., 2021).

leading 5% of commuters to discover a more-preferred alternative to their old habit. This suggests that the original behaviour was not optimal by revealed preference, nor were the magnitude and prevalence of the time savings rationalizable with search costs.[6] Larcom et al. (2017) speculate that this forced experimentation jolted commuters out of their habitual commute. Our theory provides an explanation for how sub-optimal choices can persist in stable environments, but then corrected after a shock leads to re-optimization.

2 Model

At each period $t = 1, \ldots, T$, a consumer i faces a choice among J products, indexed by $j = 1, \ldots, J$. The utility from choosing product j in period t is denoted by $u_{j,t}$, while the consumer's period-t choice is denoted by $y_t \in \{1, \ldots, J\}$. Since the model is consumer-specific, we drop the i subscript here and re-introduce it for estimation in Section 2.

We assume that the consumer implicitly forms predictions regarding the value of each product based on past utility realizations. These *reward predictions*, denoted by $r_{j,t}$ for product j in period t, evolve according to

$$r_{j,t} = \begin{cases} (1-\rho) r_{j,t-1} + \rho u_{j,t-1}, & j = y_{t-1}, \\ r_{j,t-1}, & j \neq y_{t-1}, \end{cases} \tag{1}$$

with $0 \leq \rho \leq 1$. Under (1), the consumer's reward predictions for a particular product is updated if and when that product is chosen, where the new reward

[6] In another example, after an earthquake in Chile led to a stock-out of leading beer brands, six percent of the most-frequent buyers of these brands stopped purchasing them persistently, and a substantial fraction of consumers experimented with new brands and did not switch back (Figueroa et al., 2019).

predictions is a weighted average of the previous reward prediction and the realized utility from their choice. If $\rho = 1$, then utilities are learned immediately.

Next, we assume that the consumer tracks the reliability of these reward predictions through a category-level *doubt stock*, d_t, that evolves according to

$$d_t = (1 - \lambda)d_{t-1} + |u_{j,t} - r_{j,t}| \quad \text{for } j = y_{t-1} \tag{2}$$

with $0 \leq \lambda \leq 1$. Under (2), the doubt stock is updated based on the *magnitude* of the current reward prediction error for the previously chosen product: $|u_{y_{t-1},t} - r_{y_{t-1},t}|$.[7] A large deviation between the utility of the product and the associated reward prediction thus increases the consumer's "doubt" in their prediction. As λ grows, past comparisons matter less in constructing the current doubt stock. Note that if $\rho = 1$, then the doubt stock simplifies to an accumulation of comparisons between current and past utilities, $|u_{y_{t-1},t} - u_{y_{t-1},t-1}|$.

Lastly, we assume that the consumer uses habit-based choice — in which the consumer simply repeats their choice in the previous period — if the doubt stock is below some threshold $\theta \geq 0$; otherwise, the consumer maximizes utility. Formally,

$$y_t = \begin{cases} y_{t-1}, & d_t < \theta, \\ \arg \max\limits_{j \in \{1 \dots J\}} \{u_{j,t}\}, & d_t \geq \theta. \end{cases} \tag{3}$$

By allowing the use of habit to depend on the current doubt stock (rather than lagged) ensures that current price and quality shocks can break a habit.

[7]The specification contains two departures from the original theory of Landry et al. (2021). First, we specify a single category-level doubt stock rather than a doubt stock for each product. Second, we allow current period's utility to increase doubt. These changes allow the consumer to immediately recognize that the product has changed (due to a can size change or price change) and exit habit mode in the current period. Results for the multiple doubt stock model are reported in Appendix C where we observe a small, but significant, improvement in fit. The conclusions from this model are similar.

So a consumer who is inspecting a new can, or is exposed to a new price promotion, can use this information in their immediate purchase decision. However, the consumer's ability to respond, or attend to, utility or price shocks for previously unchosen items may be limited. If the doubt stock associated with the consumer's previous choice is sufficiently low, the consumer will not actively seek information about other products.

Note that (3) reduces to a standard random utility choice model if $\theta = 0$. This suggests that the hypothesis of a 'habitual automaticity' can be directly tested through a log-likelihood comparison to a Logit model (which may or may not include a state-dependent component). In this model, all changes in observables (e.g. prices) alter the choice probabilities on all purchase occasions. By contrast, when $\theta > 0$, sufficiently small price changes for the chosen product, and all price changes for unchosen products, will not effect the choice probabilities of habitual consumers. We will elaborate on this test in Section 4.2.1.

2.1 Simulated Example

Figure 1 illustrates the model applied to a choice between two products with utilities that vary over time. In the initial periods, the consumer collects information, chooses the product with the highest utility, and forms a prediction about the utility of the chosen product for next period. Since the environment is relatively stable, prediction errors are small and the doubt stock gradually drops below threshold, with a habit formed for product 1. While in habit mode, the consumer only attends to the utility of their consumed product to assess the reliability of their utility predictions. Therefore during period 25–49 they are unaware that other products may have improved since they last made an active choice. However in period 50, the chosen product undergoes a utility shock that is immediately recognized by the consumer (e.g. a new can is introduced). The doubt stock increases, jolting the consumer out of habit mode, and back into utility maximization. The consumer continues to maximize utility for multiple periods while learning that prediction errors are small again, eventually forming a habit for product 2.

Note that, in this example, relatively more time is spent in habit mode than in maximization mode. This occurs because the choice environment (utilities) are relatively stable, with only infrequent shocks to the utility of chosen items. It is this relative stability that allows the consumer to "switch to autopilot" and avoid the associated costs of active information processing that is required to make a utility maximizing choice (e.g. assessing the prices or ingredients of all products). In effect, choice in any period is likely to be utility maximizing because choice in the previous period was.

By contrast, in an unstable environment in which the consumer was frequently confronted with prediction errors for their chosen product, our theory would predict that consumers would predominantly be in an active decision mode. This relationship between the reliability of the choice environment and habitual choice not only predicts how much a consumer will be habitual, but also when a consumer will be habitual.

3 Data

Our dataset is constructed from a household panel of purchases of 6- or 5-ounce canned tuna consumer packaged goods category, provided by AC Nielsen. We

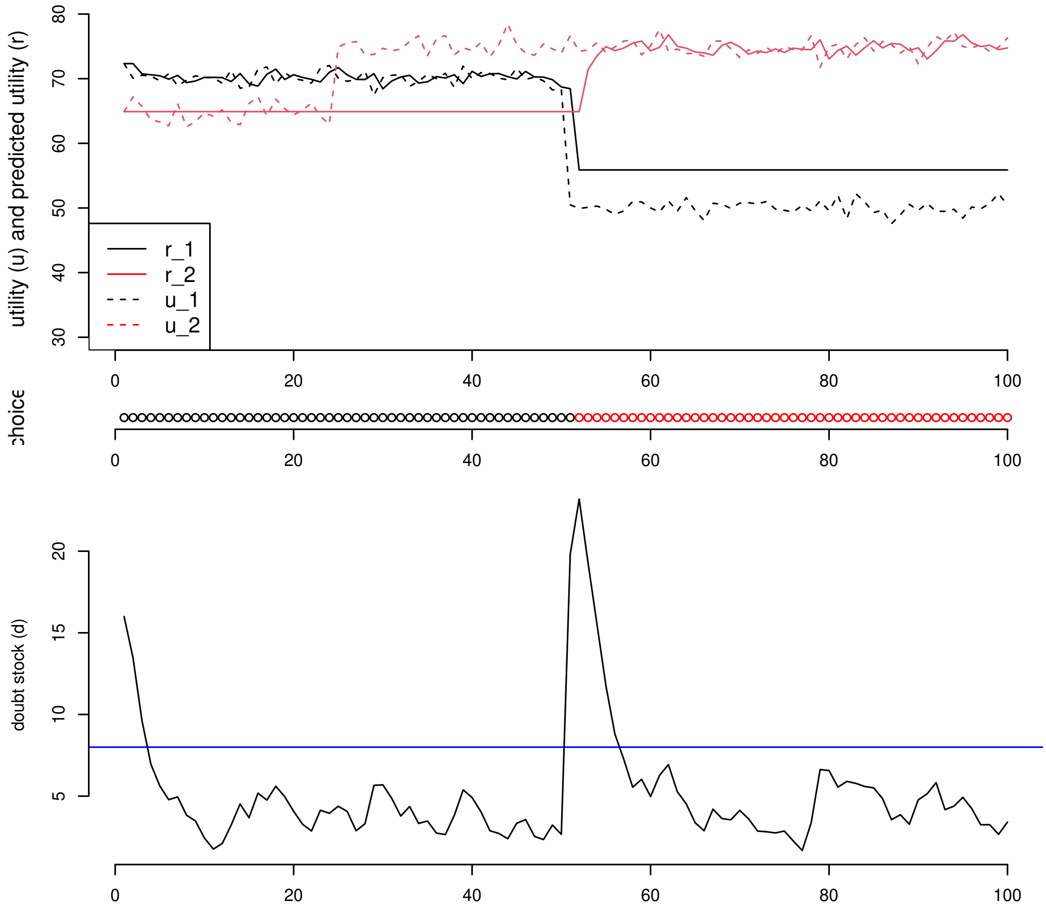

Figure 1: Simulation of utilities and doubt stock for two products. The consumer begins by maximizing utility (choosing product 1), but develops a habit as prediction errors are small. At time 25, product 2 has a utility increase, but since it is unchosen, the consumer is unaware and behaviour persists. At time 50, product 1 has a utility decrease, jolting the consumer out of their habit.

analyze the top 13 products from the three major brands in the canned tuna category, which comprise more than 70% of market share. We include only households who exclusively purchase from the chosen products and who have at least 3 purchases of canned tuna each year. Our final dataset consists of 12,524 purchase occasions from 627 households between 2006 and 2010. Weekly price data is obtained from the AC Nielsen store-level scanner panel data, in addition to product availability and whether the product had a feature or display that week.

Figure 2: Before and after packaging change for StarKist. In addition to the volume difference, the new packaging design is more modern and highlights more health-related information, such as the presence of Omega-3 fats.

In 2008, the three major brands reduced the size of canned tuna from 6 ounces to 5 ounces and changed their package design (Figure 2). Brand 1 was the first brand to begin introducing its new packaging into stores in July 2008, followed shortly by the other two brands (Figure 3). The introduction time varies not only across brands but also across stores.[8] These differences lead to natural variation in the timing of when each consumer is confronted with the product change.

[8]Specifically, 53 percent of the variation is explained by different retailers (e.g., Walmart and Safeway), while only 4 percent by geographic location (i.e., stores' DMA). This may suggest that the last contract between a retailer and a manufacturer prior to the downsizing requires the manufacturer to provide the old package up until a certain date, and the introduction date is determined by when the old contract ends and a new contract starts.

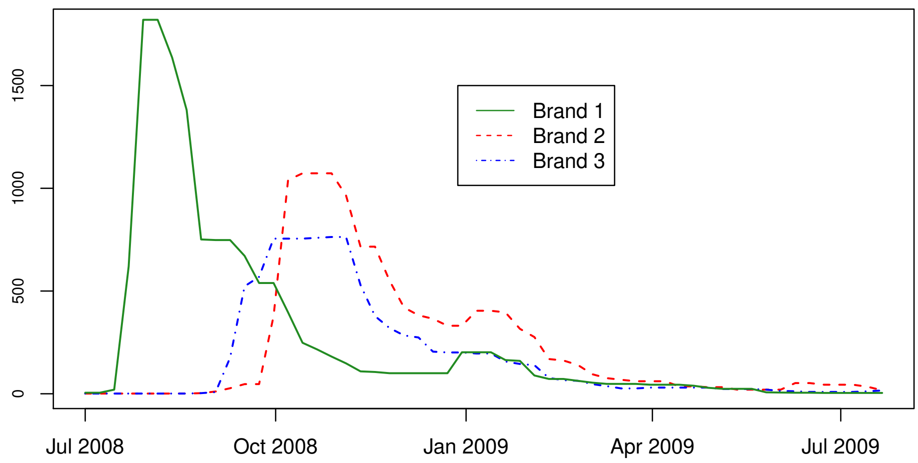

Figure 3: Number of stores that introduced new packaging of each brand during the downsizing event.

Figure 4 depicts the market shares and prices of the the three major brands during this period. During the can introduction, prices undergo a stark increase which compounds the decrease in can size for consumers, reflecting a stark instability in the choice environment facing tuna consumers. During this period, there is a noticeable re-allocation of market share that persists post-introduction. In the next section, we investigate how consumers respond to this instability in the canned tuna category, in particular how it impacts habitual choice behaviour.

4 Empirical Results

4.1 Reduced Form Analysis

We begin with a reduced-form analysis focused on the consumer's brand-switching behaviour and price sensitivity around the can change. For each consumer, we

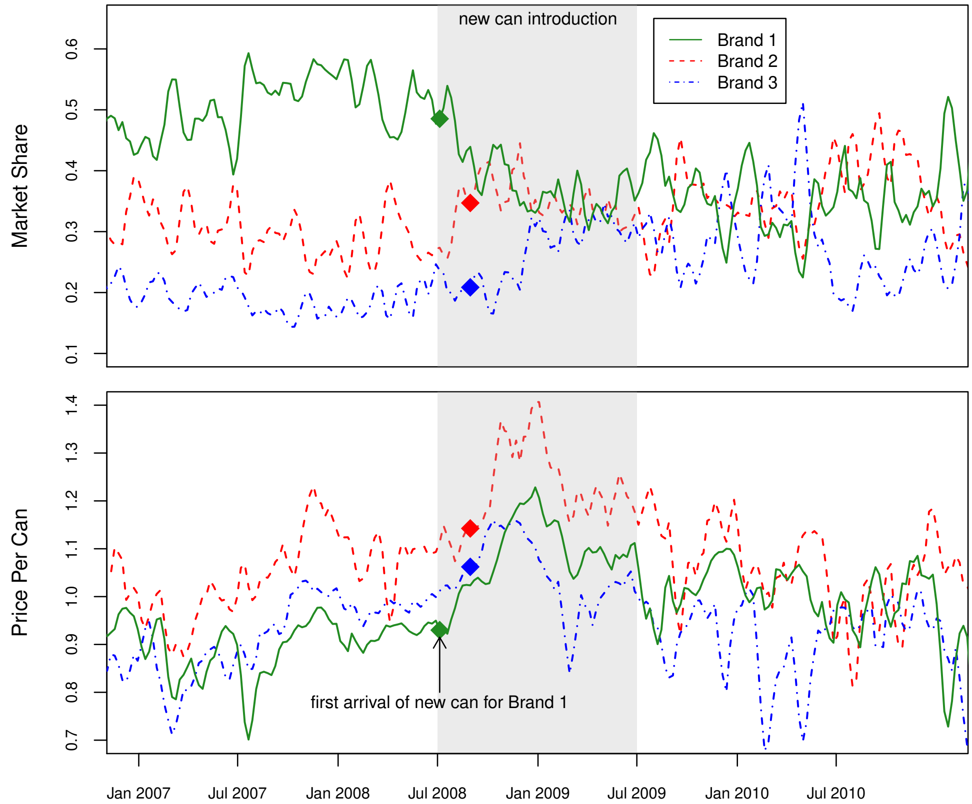

Figure 4: Market shares and prices for the three major brands

define three time windows: a period prior to the consumer's first encounter with a new can (window 1), after the first encounter but before all three brands have changed can size (window 2), and after the introduction of new can for all three brands (window 3).[9] To assess the effect of the can change in this reduced-form analysis, we categorize some consumers as habitual buyers based on their choices in window 1. Specifically, if 95% of a consumer's choices during window 1 are of the same brand, we classify the consumer as habitual.[10]

[9]The cutoffs of the three time windows can be different across consumers because stores introduce the new cans at different times and consumers shop in different stores.

[10]We report the following analysis with cutoffs of 85% and 90% in Appendix B.

To assess the degree of brand-switching, we run a mixed effect logit regression with a regressor for repeated choice, allowing for random coefficients on price and brand intercepts:

$$\begin{aligned}u_{ijt} =& \alpha_{ij} + \beta_1 \mathrm{I}(y_{i,t-1} = j) + \beta_2 \mathrm{I}(y_{i,t-1} = j) * \mathrm{habitual}_i \\ &+ \beta_3 \mathrm{I}(y_{i,t-1} = j) * \mathrm{window}_{it} + \beta_4 \mathrm{I}(y_{i,t-1} = j) * \mathrm{habitual}_i * \mathrm{window}_{it} \\ &+ \gamma_1 p_{ijt} + \gamma_2 p_{ijt} * \mathrm{habitual}_i + \gamma_3 p_{ijt} * \mathrm{window}_{it} \\ &+ \gamma_4 p_{ijt} * \mathrm{habitual}_i * \mathrm{window}_{it} + \epsilon_{ijt} \qquad (4)\end{aligned}$$

The lagged choice variable $\mathrm{I}(y_{i,t-1} = j)$ identifies whether consumer i repeated their previous choice. The dummy variable $\mathrm{habitual}_i$ indicates whether consumer i is categorized as habitual or not, while the dummy variable window_{it} indicates whether purchase occasion t is in window 1, 2 or 3 for consumer i. The average price of each brand (across UPCs for that brand) is p_{ijt}, and we allow for normally distributed coefficient correlation for price sensitivity and brand intercepts.

The estimation results of the mixed effects logit model are reported in Table 1. For non-habitual consumers, the lagged-choice parameter is positive suggesting some degree of persistence in their choices. However during the can introduction (window 2), this persistence disappears. For habitual consumers, the lagged choice parameter is much higher, not surprisingly, since they are defined to be more persistent. However, during (and after) the can introduction, we observe far less persistence in their choices. At the same time, habitual consumers become far more price sensitive. Together, these reduced-form results suggest that both brand switching and price sensitivity increased during the new can introduction, particularly for those habitual consumers whose prior purchasing behaviour was highly persistent.

	Estimate	Std. Error	
lagged choice	0.24	0.12	*
lagged choice * window 2	-0.39	0.19	*
lagged choice * window 3	-0.22	0.20	
lagged choice * habitual	4.65	0.57	***
lagged choice * habitual * window 2	-3.20	0.61	***
lagged choice * habitual * window 3	-3.58	0.67	***
price	-4.03	0.64	***
price * window 2	0.69	0.46	
price * window 3	-0.29	0.55	
price * habitual	-0.51	1.62	
price * window 2 * habitual	-3.37	1.67	*
price * window 3 * habitual	0.16	1.76	
brand 2	0.02	0.14	
brand 3	1.13	0.21	***
sd: price	1.27	0.15	***
sd: brand 2	2.35	0.23	***
sd: brand 3	4.37	0.44	***
sd: price * brand 2	0.25	0.25	
sd: price * brand 3	0.87	0.55	
sd: brand 2 * brand 3	-5.94	0.58	***

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Mixed Effects Logit Results

Given the simultaneous price increase which accompanied the can introduction, the variation of both of these variables might be leading consumers to break their habits. To assess the role of price variation, Table 2 reports an estimate of price elasticity from our sample conditioned on weeks with price discounts larger than 30%. Estimates of price elasticity appear to significantly differ depending on the size of the price change.

	Estimate	Std. Error	
lagged choice	0.32	0.08	***
lagged choice * week w/ discount > 30%	0.44	0.22	
price	-3.74	0.50	***
price * week w/ discount > 30%	-5.80	0.91	***
brand 2	-0.09	0.15	
brand 3	0.97	0.20	***
sd: price	1.55	0.16	***
sd: brand 2	2.63	0.24	***
sd: brand 3	4.68	0.49	***
sd: price * brand 2	0.27	0.24	
sd: price * brand 3	0.80	0.53	
sd: brand 2 * brand 3	-6.18	0.59	***

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Price Elasticity and Discount Level

4.2 Structural Neuro-Autopilot Results

4.2.1 Utility Specification and Choice Likelihood

In our structural analysis, we model the choice of a product (a UPC) over the time window 2006-2010 spanning the introduction of the new can. Consumer flow utility for product j is given by

$$u_{ijt} = \boldsymbol{x}'_{ijt}\boldsymbol{\beta}_i - \gamma_i p_{ijt} + \varepsilon_{ijt}, \tag{5}$$

where $\boldsymbol{x}_{ijt}$ is a vector of product characteristics with utility coefficients $\boldsymbol{\beta}_i$ (described in Table 3). These product characteristics include a random effect for each product. The price of a can is denoted p_{ijt} with a random coefficient for the marginal utility of money $\gamma_i > 0$ distributed log-normal. We assume a product-specific error ε_{ijt} is i.i.d. across time, individuals and products with CDF given by the Gumbel distribution. Additionally, in some specifications

Size	the size of the can (in ounces)
New Can Arrival	a dummy coding the first month the new can arrives in the store
Feature	whether the UPC was featured in store
Display	whether a display advertised the UPC in store
Alternative Random Effect	a random effect for each product
Brand Time Trend	a time-trend for each brand

Table 3: Observables that enter utility specification

we will include a lagged choice term in utility to construct a nested test of the state-dependent utility model:

$$u_{ijt} = \beta_1 I(y_{i,t-1} = j) + \boldsymbol{x}'_{ijt}\boldsymbol{\beta}_i - \gamma_i p_{ijt} + \varepsilon_{ijt}, \tag{6}$$

We also include a random coefficient specification for the threshold parameter θ in order to allow heterogeneity in the likelihood that a consumer enters a habit. We assume θ also follows a log-normal distribution. Our test of the habit model will therefore encompass a likelihood ratio test on these log-normal parameters.

The choice probabilities in the model are non-standard since an active choice is only made if the doubt stock is below threshold. Since the doubt stock is a function of past choices — past draws of ε_{ijt} for chosen products — standard simulation approaches from the unbounded Gumbel distribution can make the log-likelihood discontinuous in some parameters. Therefore we use an importance-sampling technique similar to the procedure used to estimate dynamic Tobit models (e.g., Lee 1999). To guide intuition, if we observe a consumer switch products, we know the error draw for the previously chosen product must have arisen within some known range and we can draw this error from a truncated Type 1 extreme value distribution. Importantly, if the threshold is zero (therefore no choices are made in habit mode) then these choice probabilities collapse to the standard logit probability. Full details of the expressions can be found in Appendix A.

To simulate the likelihood, we initialize a consumer's doubt stock given a history of price changes faced by that consumer and simulate 500 paths of error draws to construct simulated doubt stocks.[11] The average likelihood is then maximized with with an adaptive algorithm followed by gradiant descent. Standard errors are calculated from the outer-product gradiant.

4.2.2 Neuro-Autopilot Estimates

The results from our model estimation are presented in Table 4. In columns (1) and (2) we present the standard random-coefficients Logit with and without a lagged choice variable. Columns (4) and (5) present our neuro-autopilot model, again with and without lagged choice.

In both specification (4) and (5), we find that the density for the habit threshold (θ) places substantial weight above zero and we can definitively reject a baseline random utility model (1) or (2) in which consumer are actively comparing utilities on every purchase occasion (LR test of (4) vs. (1), $p < 10^{-8}$; LR test of (5) vs./ (2), $p < 10^{-8}$). While the coefficient on lagged choice in (2) is positive and significant, we find this model performs worse than even our baseline autopilot model (ΔBIC (4) vs. (2): 461). When we include a lagged

[11] To be more specific, at each evaluation of the likelihood an initial choice is simulated given current parameter values and price histories to form the initial doubt stock. This consumer-specific initialization therefore does not require an assumption that all consumers are either in or out of habit.

Parameter	(1) Base Logit	(2) Logit w Lag Choice	(3) Mix Logit w Rep Choice	(4) Base Autopilot	(5) Autopilot w Lag Choice
Prediction Depr (ρ)	-	-	-	0.48	0.11
				(0.04)	(0.02)
Doubt Depr (λ)	-	-	-	0.98	0.70
				(0.02)	(0.03)
Threshold (θ) mean	-	-	-	0.68	0.87
				(0.03)	(0.14)
Threshold (θ) sd	-	-	-	1.41	2.78
				(0.44)	(1.63)
Price (γ) mean	-5.41	-5.38	-5.86	-6.64	-6.67
	(0.46)	(0.58)	(0.63)	(0.99)	(0.95)
Price (γ) sd	5.14	4.99	5.51	6.69	6.41
	(0.42)	(0.53)	(0.53)	(0.76)	(0.71)
Size (ounces)	0.25	0.24	0.29	0.08	0.41
	(0.09)	(0.09)	(0.11)	(0.07)	(0.07)
New Can Arrival month	0.41	0.47	0.37	0.46	0.57
	0.29	(0.29)	(0.33)	(0.28)	(0.25)
Lag Choice	-	0.88	-	-	0.53
		(0.02)			(0.03)
Switch Prob	-	-	0.87	-	-
			(0.00)		
Log- Likelihood	-14501	-14082	-14145	-13832	-13815
BIC	29305	28475	28601	28014	27988

Table 4: Model Estimates

choice term in our autopilot model (5), we find that the habit threshold is larger, and the magnitude of the lagged choice coefficient decreases by 40%. Together, these results suggest that there may be two channels in which choice persistence arises in consumer purchases, a primary channel through a "naive" habitual autopilot mode in which choices are automatically repeated, and a secondary channel through preference complementarity.

To assess the model's predictions about *when* consumer are in habit mode, in column (3) we present a "mixed" version of the logit model in which consumer

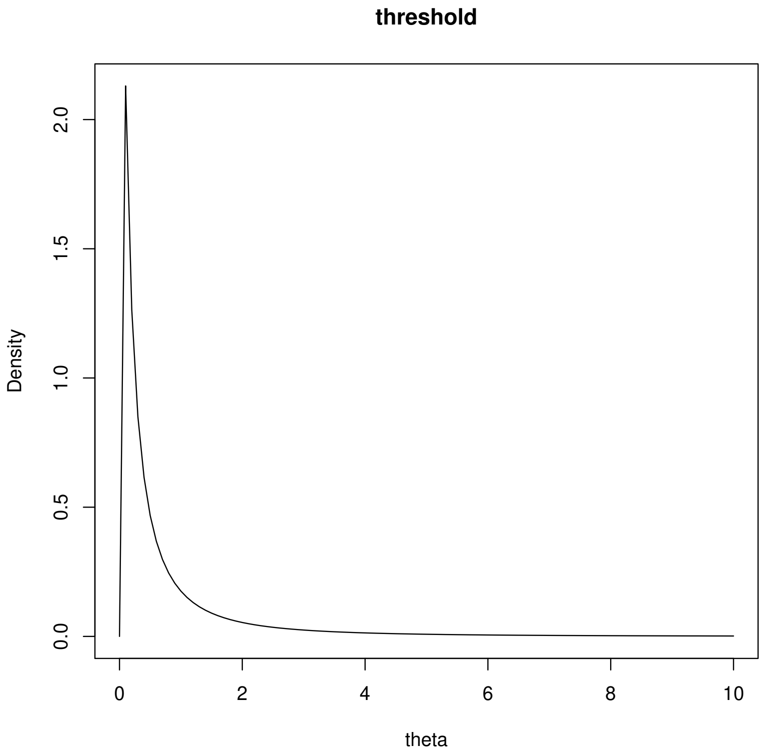

Figure 5: Estimated density of θ

are in habit vs. random utility maximization with some probability that is independent over purchase occasions. This lies in contrast with our autopilot model which specifies when habit mode should arise: when observables like prices and can size are relatively stable. The improved performance of the autopilot model (4) therefore reflects the timing of habit mode, as specified by the doubt stock relative to threshold, coinciding with periods of relative stability in prices and product characteristics.

Conversely, it is the variation in these observables around the time of the new can introduction that leads to re-optimization. One feature of our habit model is that we can calculate the expected proportion of consumers in habit mode at any period, conditional on observables. We present these calculations for three models in Figure 6. In our preferred specification (5), we calculate that 12% of canned tuna choices are made in habit mode. During the can introduction, this percentage drops to 10%.

The parameters governing the reward prediction and doubt stock processes are also of interest. Since $\rho = 0.11$ is significantly lower than 1, utility predictions aggregate over many periods rather than just the previous period's utility. The depreciation of the doubt stock ($\lambda = 0.70$)is also less than 1, suggesting reward reliability is also tracked over multiple periods. Therefore exiting a habit involves more than just a direct comparison of utility with the previous period, rather it involves an aggregation of surprising utility shocks over many periods (with some decay). This allows the agent to strike a balance between allowing a single event vs. a sequence of "surprising-enough" events to lead to a re-optimization; yielding a trade-off between exploration and exploitation that can balance the cost of collecting utility information on every trial.[12] This

[12]Interestingly, the form of trade-off employed here differs from the standard "ϵ-greedy" approach

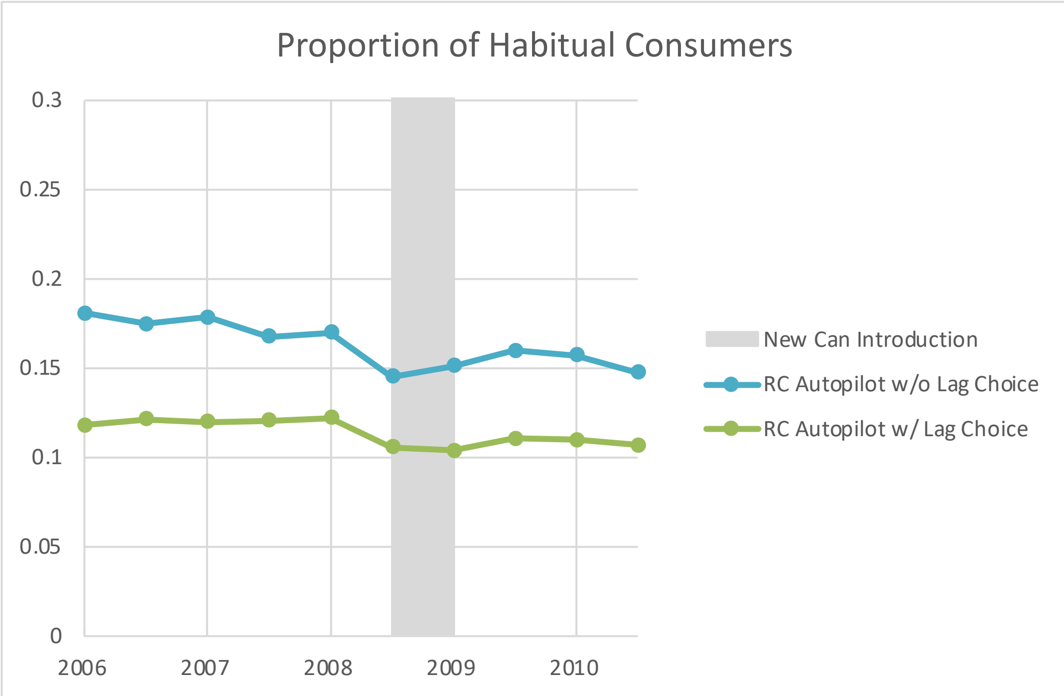

Figure 6: Proportion of Habitual Consumers

implies that when consumer's exit habit mode because of a price/reward shock, they don't immediately re-enter habit mode upon their next purchase. Rather, the doubt stock will persist above threshold for a number of periods and the consumer will actively compare utility information before settling into a new habit.

5 Conclusion

We have proposed a model of consumer choice which includes a habitual decision-making mode in addition to preference-based choice. The arbitration of these

in machine learning, in which exploration is simply a random perturbation of behaviour and it is exploitation that results from a maximization process. In our model, the agent explores via a maximization process when the environment is unstable, and is persistent it is not.

two systems is drawn from previous neuroeconomic research on how the human brain tracks the reliability of reward signals. We estimated this model on consumer purchases of canned tuna, and we find that the hypothesis of an active random-utility maximizer on every purchase occasion can be summarily rejected in favour of a habitual consumer who automates their choices when the choice environment is stable ($\sim 12\%$ of consumption occasions). Moreover, a majority of the choice persistence in our data is due to habitual automation rather than direct effects on utility. Though we do still find a role for past choices to enter utility, the contribution of this term in explaining choice persistence is smaller, suggesting its role may be over-stated.

References

Abaluck, J. and Adams-Prassl, A. (2021). What do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses. *The Quarterly Journal of Economics*, 136(3):1611–1663.

Adams, C. D. and Dickinson, A. (1980). Instrumental Responding following Reinforcer Devaluation. *Quarterly Journal of Experimental Psychology*, 33(2b):109–121.

Auld, M. and Grootendorst, P. (2004). An empirical analysis of milk addiction. *Journal of Health Economics*, 23(6):1117–1133.

Becker, G. S. (1996). *Accounting for Tastes.*

Becker, G. S. and Murphy, K. M. (1988). A Theory of Rational Addiction. *Journal of Political Economy*, 96(4):675–700.

Bronnenberg, B. J., Dubé, J.-P. H., and Gentzkow, M. (2012). The Evolution of Brand Preferences: Evidence from Consumer Migration. *American Economic Review*, 102(6):2472–2508.

Cerigioni, F. (2021). Dual Decision Processes: Retrieving Preferences When Some Choices Are Automatic. *Journal of Political Economy*, 129(6):1667–1704.

Crawford, I. (2010). Habits Revealed. *Review of Economic Studies*, 77(4):1382–1402.

Dubé, J., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445.

Figueroa, C., Musalem, A., and Noton, C. (2019). Earthquakes and Brand Loyalty: Beyond the Short-Term Effects of Stockouts. *SSRN Electronic Journal*.

Giuliano, P. and Nunn, N. (2021). Understanding Cultural Persistence and Change. *The Review of Economic Studies*, 88(4):1541–1581.

Heckman, J. J. (1981). Statistical models for discrete panel data. In Manski, C. F. and McFadden, D., editors, *Structural Analysis of Discrete Data with Applications*.

Keane, M. P. (1997). Modeling Heterogeneity and State Dependence in Consumer Choice Behavior. *Journal of Business & Economic Statistics*, 15(3):310.

Kong, X., Dube, J.-P. H., and Daljord, O. (2022). Non-parametric Estimation of Habitual Brand Loyalty. *SSRN Electronic Journal*.

Landry, Webb, and Camerer (2021). A Neural Autopilot Theory of Habit. *SSRN Working Paper*.

Landry, P. (2019). Bad Habits and the Endogenous Timing of Urges. *The Review of Economic Studies*, 86(2):785–806.

Laporte, A., Dass, A. R., and Ferguson, B. S. (2017). Is the Rational Addiction model inherently impossible to estimate? *Journal of Health Economics*, 54:161–175.

Larcom, S., Rauch, F., and Willems, T. (2017). The Benefits of Forced Experimentation: Striking Evidence from the London Underground Network*. *The Quarterly Journal of Economics*, 132(4):2019–2055.

Lee, S. W., Shimojo, S., and O'Doherty, J. P. (2014). Neural Computations Underlying Arbitration between Model-Based and Model-free Learning. *Neuron*, 81(3):687 – 699.

Manzini, P. and Mariotti, M. (2014). Stochastic Choice and Consideration Sets. *Econometrica*, 82(3):1153–1176.

Masatlioglu, Y., Nakajima, D., and Ozbay, E. Y. (2012). Revealed Attention. *American Economic Review*, 102(5):2183–2205.

Matysková, L., Rogers, B., Steiner, J., and Sun, K.-K. (2020). Habits as adaptations: An experimental study. *Games and Economic Behavior*, 122:391–406.

Olekalns, N. and Bardsley, P. (1996). Rational Addiction to Caffeine: An Analysis of Coffee Consumption. *Journal of Political Economy*, 104(5):1100–1104.

Pollak, R. A. (1970). Habit Formation and Dynamic Demand Functions. *Journal of Political Economy*, 78(4, Part 1):745–763.

Pool, E. R., Gera, R., Fransen, A., Perez, O. D., Cremer, A., Aleksic, M., Tanwisuth, S., Quail, S., Ceceli, A. O., Manfredi, D. A., Nave, G., Tricomi, E., Balleine, B., Schonberg, T., Schwabe, L., and O'Doherty, J. P. (2021). Determining the effects of training duration on the behavioral expression of habitual control in humans: a multilaboratory investigation. *Learning & memory (Cold Spring Harbor, N.Y.)*, 29(1):16–28.

Ryder, H. E. and Heal, G. M. (1973). Optimal Growth with Intertemporally Dependent Preferences. *The Review of Economic Studies*, 40(1):1.

Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306):1593 – 1599.

Seetharaman, P. B. (2004). Modeling Multiple Sources of State Dependence in Random Utility Models: A Distributed Lag Approach. *Marketing Science*, 23(2):263–271.

Shum, M. (2004). Does Advertising Overcome Brand Loyalty? Evidence from the Breakfast-Cereals Market. *Journal of Economics & Management Strategy*, 13(2):241–272.

Siegel, S., Hinson, R. E., Krank, M. D., and McCully, J. (1982). Heroin "Overdose" Death: Contribution of Drug-Associated Environmental Cues. *Science*, 216(4544):436–437.

Steiner, J., Stewart, C., and Matějka, F. (2017). Rational Inattention Dynamics: Inertia and Delay in Decision-Making. *Econometrica*, 85(2):521–553.

Thomadsen, R. and Seetharaman, P. B. S. (2018). In Verplanken, B., editor, *The Psychology of Habit*, pages 131–150.

Wood, W., Mazar, A., and Neal, D. T. (2021). Habits and Goals in Human Behavior: Separate but Interacting Systems. *Perspectives on Psychological Science*, page 174569162199422.

Wood, W. and Neal, D. T. (2009). The habitual consumer. *Journal of Consumer Psychology*, 19(4):579–592.

Appendix A Choice Probability

Denote consumer i's flow utility at time t for product j as

$$u_{ijt} = \overline{u}_{ijt} + \varepsilon_{ijt}.$$

For simplicity, denote the consumer's last period choice as k. The consumer will be in habit mode when

$$\begin{aligned}
& (1-\lambda)d_{i,t-1} + |u_{ikt} - r_{ikt}| < \theta \\
\iff \quad & -\theta + (1-\lambda)d_{i,t-1} < \overline{u}_{ikt} + \varepsilon_{ikt} - r_{ikt} < \theta - (1-\lambda)d_{i,t-1} \\
\iff \quad & r_{ikt} - \overline{u}_{ikt} - \theta + (1-\lambda)d_{i,t-1} < \varepsilon_{ikt} < r_{ikt} - \overline{u}_{ikt} + \theta - (1-\lambda)d_{i,t-1} \\
\iff \quad & a < \varepsilon_{ikt} < b.
\end{aligned}$$

Otherwise, the consumer will make an active choice.

A.1 Choice probability if no switch is observed

First, we consider the likelihood that a choice is repeated from period $t-1$ to t. If the consumer does not switch, there are two possibilities: (i) the individual is in habit mode ($d_{ijt} < \theta$), or (ii) the consumer is not in habit model ($d_{ijt} \geq \theta$) the previously chosen product has the highest utility in period t. To simplify notation, we drop the i and t subscripts and denote the previously chosen product as k.

The first case arises only if ε_{ikt} are in some range $[a, b]$.

Case i: Here, we need to compute the probability that $a < \varepsilon_k < b$. Noting that the logit CDF is $F(\varepsilon) = e^{-e^{-\varepsilon}}$, this probability is simply:

$$F(b) - F(a) = e^{-e^{-b}} - e^{-e^{-a}}.$$

Case ii: If the consumer makes an active, choice, her ε_k draw is outside the $[a, b]$ interval and it is above the utilities of all other products: $\varepsilon_k > \max_{j \neq k}\{\overline{u}_j - \overline{u}_k + \varepsilon_j\}$. To compute the choice probability here we will use the usual trick to compute the logit choice probabilities. Suppose that we know ε_k. The probability product k is chosen will then be

$$\prod_{j \neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k),$$

and the overall choice probability will be

$$\int_{-\infty}^{a} \prod_{j \neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k + \int_{b}^{\infty} \prod_{j \neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k.$$

Let's consider computing the first integral. We can write this as

$$\begin{aligned}
\int_{-\infty}^{a} \prod_{j \neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k &= \int_{-\infty}^{a} \prod_{j \neq k} \exp(-\exp(-(\overline{u}_k - \overline{u}_j + \varepsilon_k))) \exp(-\exp -\varepsilon_k)) \exp(-\varepsilon_k \\
&= \int_{-\infty}^{a} \prod_{j} \exp(-\exp(-(\overline{u}_k - \overline{u}_j + \varepsilon_k))) \exp(-\varepsilon_k) d\varepsilon_k \\
&= \int_{-\infty}^{a} \exp\left(-\sum_j \exp(-(\overline{u}_k - \overline{u}_j + \varepsilon_k))\right) \exp(-\varepsilon_k) d\varepsilon_k \\
&= \int_{-\infty}^{a} \exp\left(-\exp(-\varepsilon_k)\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right) \exp(-\varepsilon_k) d\varepsilon_k.
\end{aligned}$$

Now we can make a change of variables, we will set $t = -\exp(-\varepsilon_k)$, $dt = \exp(-\varepsilon_k)d\varepsilon_k$, where the bounds on t are $(-\infty, -e^{-a})$. Now we rewrite the integral as

$$\int_{-\infty}^{a} \prod_{j\neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k = \int_{-\infty}^{-e^{-a}} \exp\left(t\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right) dt.$$

So the choice probability will be

$$\begin{aligned} \int_{-\infty}^{a} \prod_{j\neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k &= \left.\frac{\exp\left(t\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))}\right]_{-\infty}^{-e^{-a}} \\ &= \frac{\exp\left(-\exp(-a)\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))}. \end{aligned}$$

Following the same logic, the second integral should then be

$$\begin{aligned} \int_{b}^{\infty} \prod_{j\neq k} F(\overline{u}_k - \overline{u}_j + \varepsilon_k) f(\varepsilon_k) d\varepsilon_k &= \left.\frac{\exp\left(t\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))}\right]_{-e^{-b}}^{0} \\ &= \frac{1 - \exp\left(-\exp(-b)\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))}. \end{aligned}$$

Hence, the overall choice likelihood in this case will be:

$$
\begin{aligned}
P(\text{don't switch from } k) \quad = \quad & e^{-e^{-b}} - e^{-e^{-a}} \\
& + \frac{\exp\left(-\exp(-a)\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))} \\
& + \frac{1 - \exp\left(-\exp(-b)\left[\sum_j \exp(-(\overline{u}_k - \overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_k - \overline{u}_j))}.
\end{aligned}
$$

Note that if $a = b$, this expression collapses to the standard logit probability.

A.2 Choice probability if a switch is observed

Suppose that the last chosen brand is brand k, and the consumer switches to brand $l \neq k$. Then we know an active choice was made. There are two cases we should consider: (i) $\varepsilon_k < a$ and (ii) $\varepsilon_k > b$. The choice probability will be the sum of the probabilities in both cases.

Case i: Again, we will treat ε_l as known. The main difference in the conditions on the other error terms is that ϵ_k must be below both $\overline{u}_l - \overline{u}_k + \varepsilon_l$ and a. In this case the probability that product l is chosen conditional on its error is

$$
\left(\prod_{j \notin \{l,k\}} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) F(\min\{\overline{u}_l - \overline{u}_k + \varepsilon_l, a\}).
$$

Note that the choice probability formula now depends on whether $\overline{u}_l - \overline{u}_k + \varepsilon_l < a$ or not. So we can write it as the sum of two integrals:

$$
\int_{-\infty}^{a-(\overline{u}_l - \overline{u}_k)} \left(\prod_{j \neq l} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) f(\varepsilon_l) d\varepsilon_l + \int_{a-(\overline{u}_l - \overline{u}_k)}^{\infty} \left(\prod_{j \notin \{l,k\}} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) F(a) f(\varepsilon_l) d\varepsilon_l.
$$

The derivations of the integrals above can be done using similar math to the case where we don't see a switch. In particular, we should get

$$\int_{-\infty}^{a-(\overline{u}_l-\overline{u}_k)} \left(\prod_{j\neq l} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) f(\varepsilon_l) d\varepsilon_l = \frac{\exp\left(-\exp(-(a-(\overline{u}_l-\overline{u}_k)))\left[\sum_j \exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_l-\overline{u}_j))},$$

and

$$\begin{aligned}
& \int_{a-(\overline{u}_l-\overline{u}_k)}^{\infty} \left(\prod_{j\notin\{l,k\}} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) F(a) f(\varepsilon_l) d\varepsilon_l \\
= \quad & F(a)\frac{1-\exp\left(-\exp(-(a-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j\neq k} \exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j\neq k} \exp(-(\overline{u}_l-\overline{u}_j))}.
\end{aligned}$$

Case ii: Treating ε_l as known, the difference here is we have both upper and lower bounds on ϵ_k. So the choice probability conditional on ε_l will be

$$\left(\prod_{j\notin\{l,k\}} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right)(F(\overline{u}_l-\overline{u}_k+\varepsilon_l)-F(b)) = \left(\prod_{j\neq l} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) - \left(\prod_{j\notin\{l,k\}} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) F$$

as long as $\overline{u}_l - \overline{u}_k + \varepsilon_l > b$, and 0 otherwise. The choice probability here will be the integral over ε_l of the terms above from $b-(\overline{u}_l-\overline{u}_k)$ to ∞. Considering the terms separately, the first term will be

$$\int_{b-(\overline{u}_l-\overline{u}_k)}^{\infty} \left(\prod_{j\neq l} F(\overline{u}_l - \overline{u}_j + \varepsilon_l)\right) f(\varepsilon_l) d\varepsilon_l.$$

We know how to evaluate this integral from the derivations above, it will be

$$\frac{1-\exp\left(-\exp(-(b-(\overline{u}_l-\overline{u}_k)))\left[\sum_j \exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_j \exp(-(\overline{u}_l-\overline{u}_j))}.$$

The second choice probability will be very similar:

$$F(b)\frac{1-\exp\left(-\exp(-(b-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))}.$$

Hence, the overall probability of a switch will be

$$\begin{aligned}
Prob(\text{switch from } k \text{ to } l) \quad = \quad & \frac{\exp\left(-\exp(-(a-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j}\exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j}\exp(-(\overline{u}_l-\overline{u}_j))} \\
& +e^{-e^{-a}}\frac{1-\exp\left(-\exp(-(a-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))} \\
& +\frac{1-\exp\left(-\exp(-(b-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j}\exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j}\exp(-(\overline{u}_l-\overline{u}_j))} \\
& -e^{-e^{-b}}\frac{1-\exp\left(-\exp(-(b-(\overline{u}_l-\overline{u}_k)))\left[\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))\right]\right)}{\sum_{j\neq k}\exp(-(\overline{u}_l-\overline{u}_j))}
\end{aligned}$$

Again if $a = b$, this expression collapses to the standard logit probability.

Appendix B More Reduced Form Results

In this appendix, we conduct robustness tests for the reduced form results in Table 1 using different definitions of habitual buyers. Previously, we classify a consumer as a habitual buyer if 95% of a consumer's choices during window 1 are of the same brand. In Table 5, we report the same analysis results but with cutoff of 85% (column 1) and 90% (column 2) when defining habitual buyers. The estimates here are of similar magnitudes and significance as in Table 1.

	(1)		(2)	
lagged choice	0.19		0.17	
	(0.12)		(0.12)	
lagged choice * window 2	-0.30		-0.24	
	(0.19)		(0.19)	
lagged choice * window 3	-0.25		-0.17	
	(0.21)		(0.21)	
lagged choice * habitual	4.12	***	3.54	***
	(0.46)		(0.37)	
lagged choice * habitual * window 2	-2.89	***	-2.28	***
	(0.51)		(0.43)	
lagged choice * habitual * window 3	-2.96	***	-2.40	***
	(0.56)		(0.48)	
price	-4.17	***	-2.66	***
	(0.65)		(0.64)	
price * window 2	0.80		0.78	
	(0.47)		(0.47)	
price * window 3	-0.21		-0.27	
	(0.56)		(0.57)	
price * habitual	-0.51		-1.81	
	(1.29)		(1.17)	
price * window 2 * habitual	-3.21	*	-3.02	*
	(1.37)		(1.24)	
price * window 3 * habitual	0.12		0.28	
	(1.43)		(1.27)	
brand 2	0.00		0.08	
	(0.15)		(0.14)	
brand 3	1.13	***	1.10	***
	(0.21)		(0.19)	
sd: price	1.31	***	1.28	***
	(0.15)		(0.14)	
sd: brand 2	2.33	***	2.43	***
	(0.23)		(0.28)	
sd: brand 3	4.38	***	3.59	***
	(0.45)		(0.40)	
sd: price * brand 2	0.28		0.68	**
	(0.24)		(0.25)	
sd: price * brand 3	0.79		-1.16	*
	(0.53)		(0.58)	
sd: brand 2 * brand 3	-5.84	***	-7.41	***
	(0.57)		(0.81)	

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Mixed Effects Logit w Different Definitions for Habitual Buyers

Appendix C Neuro-Autopilot Results for Multiple Doubt Stocks

The section reports results from a model with multiple doubt stocks as originally specificed in Landry et al. (2021). Specifically, there is a doubt stock formed for each product,

$$d_{j,t} = \begin{cases} (1-\lambda)d_{j,t-1} + |u_{j,t} - r_{j,t}|, & \text{if } y_{t-1} = j \\ (1-\lambda)d_{j,t-1} + \alpha, & \text{otherwise} \end{cases} \tag{7}$$

with $0 \leq \lambda \leq 1$. Under (7), the doubt stock associated with the consumer's choice in the previous period is updated based on the *magnitude* of the current reward prediction error: $|u_{j,t} - r_{j,t}|$. For products that were not chosen in the previous period, the doubt stock is instead updated with the addition of the constant $\alpha > 0$. Thus, the consumer's level of doubt in a reward prediction may also also increase as a result of *not* choosing the associated product, as this reward prediction has not recently been "tested" against utilities.

Behaviour is then determined by a comparison of the doubt stock for the previously chosen product

$$y_t = \begin{cases} y_{t-1}, & d_{y_{t-1},t} < \theta, \\ \arg\max_{j \in \{1 \ldots J\}} \{u_{jt}\}, & d_{y_{t-1},t} \geq \theta. \end{cases} \tag{8}$$

Results for this model are reported in Table 6. In columns (1) and (2) we present the standard random-coefficients Logit with and without a lagged choice variable. Columns (4) and (5) present our neuro-autopilot model, again with and without lagged choice.

Parameter	(1) Base Logit	(2) Logit w Lag Choice	(3) Mix Logit w Rep Choice	(4) Base Autopilot	(5) Autopilot w Lag Choice
Prediction Depr (ρ)	-	-	-	0.42	0.17
				(0.03)	(0.02)
Doubt Depr (λ)	-	-	-	0.55	0.34
				(0.03)	(0.01)
Doubt Incr (α)	-	-	-	2.80	2.70
				(0.22)	(0.18)
Threshold (θ) mean	-	-	-	1.59	2.06
				(0.09)	(0.15)
Threshold (θ) sd	-	-	-	0.92	2.20
				(0.11)	(0.61)
Price (γ) mean	-5.41	-5.38	-5.86	-6.09	-6.22
	(0.46)	(0.58)	(0.63)	(0.76)	(0.89)
Price (γ) sd	5.14	4.99	5.51	5.75	6.07
	(0.42)	(0.53)	(0.53)	(0.62)	(0.72)
Size (ounces)	0.25	0.24	0.29	0.12	0.20
	(0.09)	(0.09)	(0.11)	(0.06)	(0.08)
New Can Arrival month	0.41	0.47	0.37	0.29	0.33
	0.29	(0.29)	(0.33)	(0.25)	(0.27)
Lag Choice	-	0.88	-	-	0.62
		(0.02)			(0.03)
Switch Prob	-	-	0.87	-	-
			(0.00)		
BIC	29305	28475	28601	28073	27976

Table 6: Model Estimates

Results are consistent with the single doubt stock model previously reported. In both specification (4) and (5), we can definitively reject a random utility model (1) or state-dependent utility (2) in which consumer are actively comparing utilities on every purchase occasion (LR test of (4) vs. (1), $p < 10^{-8}$; LR test of (5) vs./ (2), $p < 10^{-8}$). While the coefficient on lagged choice in (2) is positive and significant, we find this model performs worse thanthe baseline autopilot model (ΔBIC (4) vs. (2): 402). When we include a lagged choice term in our autopilot model (5), we find that the habit threshold is larger, and the magnitude of the lagged choice coefficient decreases by roughly a third.

Overall, the multiple doubt stock models provide a small but significant improvement in fit compared to their companion single doubt stock specifications, but no qualitiative differences. In our preferred specification (5), we calculate that 11% of canned tuna choices are made in habit mode. During the can introduction, this percentage drops to 9%.

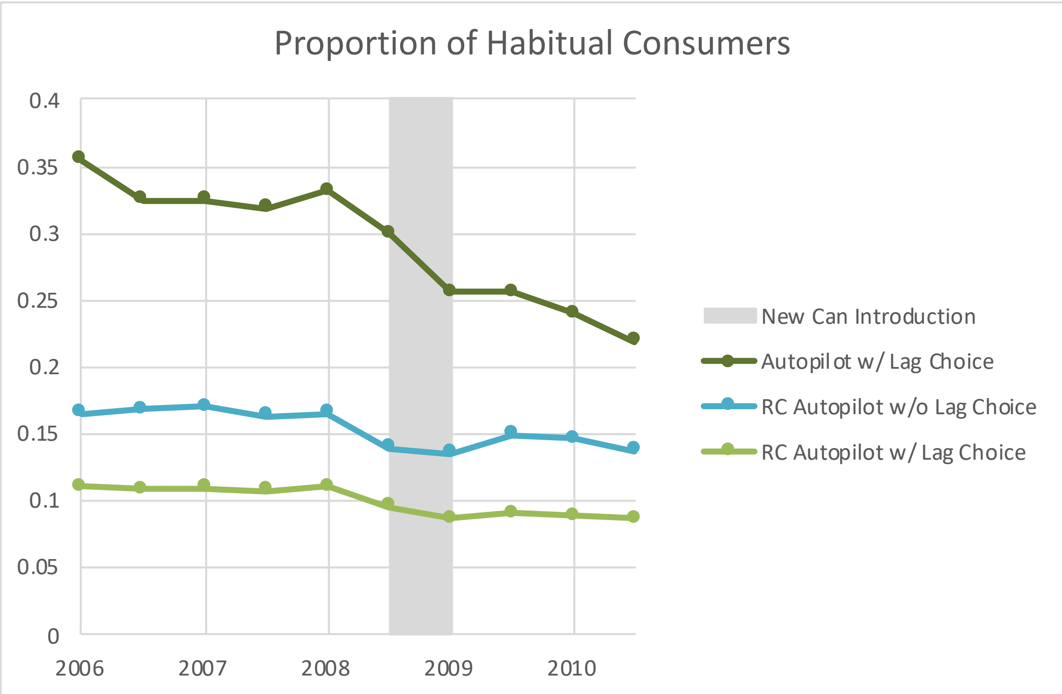

Figure 7: Proportion of Habitual Consumers

The main difference between the two specifications is that the multiple doubt stock model contains an extra parameter, α which determines the steady-state of the doubt stock ($\frac{\alpha}{\lambda}$) when an item is unchosen. When compared to θ this determines the probability that a consumer immediately re-enters habit mode after exiting. We estimate $\frac{\alpha}{\lambda} = 7.94$, which corresponds to the upper tail of the density for θ. This implies that when consumer's exit habit mode because of a price/reward shock, they don't immediately re-enter habit mode upon their next purchase.

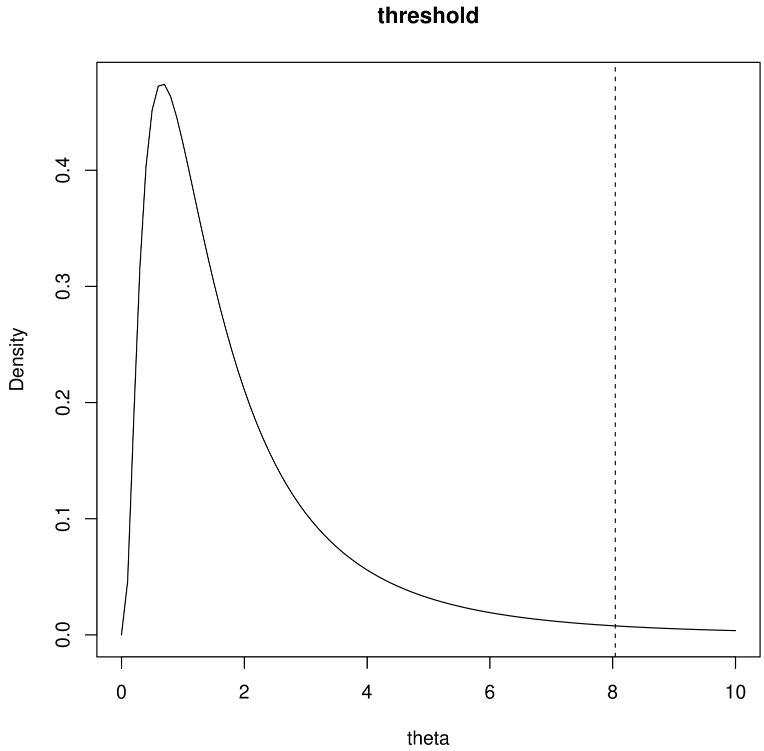

Figure 8: Estimated density of θ