A neural autopilot theory of habit

Peter Landry, Ryan Webb, Colin Camerer

November 7, 2021
Prepared for "The State of Mind in Economics" (A Kirman and M Teschi, Eds.)

I. Introduction

This chapter presents a neuroeconomic theory (i.e. a simplified computational model) of habitual choice. While grounded in neural mechanisms of habit, the theory accounts for some basic facts about consumer behavior (with interesting new predictions).

To understand how insights from neuroscience can be used in economics more generally, we define economics as **the analysis of consequential human choices under scarcity and institutional constraints, including their causes and impacts on human welfare**. This definition is crafted to constrain the ingredients of an economic theory, but still permits a lot of variation. For example, even though the definition refers to *humans,* human biology has much in common with the biology of other animals (especially other primates). Evidence from other animals – e.g. through species-general reactions to visual designs, mental states, group dynamics, or status incentives – can therefore help inform economic theories. Also note, this definition does not privilege a particular type of data or methodology, in contrast to assertions such as "revealed preference earns such a central role in economics because this [choice] is the form of evidence that is available to economists" (Gul and Pesendorfer, 2008).

a. Three levels of understanding

A guiding methodological framework for neuroeconomics is the three-level analysis introduced by the vision scientist David Marr (1982). Marr was interested in how we "see" the world around us, but was frustrated by how little scientific attention was paid to integrating the mechanistic details of vision (i.e. the anatomy and physiology of the visual system) with a description of their function (i.e. how we might program a computer to see). Marr wrote: "trying to understand perception by understanding neurons is like trying to understand a bird's flight by studying only feathers. It just cannot be done" (p. 27).

Marr described three levels at which behavior can be understood. The highest level, **functionality** (our preferred term, though Marr referred to this as "computation"), addresses *why* the brain is organized in a particular way to solve some problem facing the organism. The middle level, **algorithm**, explains *what* is computed to serve this function. The lowest level, **implementation**, describes *how* this computation is executed through biological processes. Returning to the example of bird flight, Krakauer et al. (2017) suggest that flight — useful for evading predators, nesting, and migration — is the highest-level function of the bird's biological system; the middle-level "algorithm" to achieve flight can be described by the equations that govern the flapping of wings; the lowest-level "implementation" then refers to the precise biological mechanisms used by the bird to flap its wings.

The key premise of this three-level framework is that behavior can only be well understood if it is well understood at all three levels. This suggests that theoretical models of behavior should be compatible with evidence at all three levels; it also suggests that the representation of behavior at each level should be informed by the representations at other levels.[1] In our view, Marr's three-level framework can and should be applied to the social sciences, though microeconomic theories of individual choice often neglect the lowest levels of algorithm and implementation, while focusing only on describing behavior at the functional level (if only partially).

[1] As an illustration, hummingbirds do not fly solely as a mode of transport. They also hover to reach otherwise inaccessible nectar sources (functionality). To serve this alternate function, hummingbirds beat their wings far more rapidly than other birds while generating lift on both up and down beats (algorithm). Meanwhile, hummingbirds' rapid wing-beating is made possible by disproportionately large pectoralis muscles (implementation).

For example, consider standard choice theories featuring agents who make choices that maximize their subjective utility. These theories only partially address — through behavioral axioms — the functional question of why agents maximize utility as they do not incoporate relevant constraints on attention and cognition. Given two organisms, the one that most efficiently allocates scarce cognitive resources in processing utility-relevant information will have an advantage — even if it might make some mistakes. And of course standard utility theories are silent on the algorithms that people might use to implement (or approximate) utility maximization, and their actual physical implementation subject to biophysical constraints.

A few economic theories have taken on the challenge of integrating Marr's three levels, using different methods. One path derives the types of utility functions that would emerge under selective pressure in particular choice environments (e.g. Robson, 2002, Rayo & Becker, 2007).[2] As an example, it is commonly held that risks are valued by weighting the utility of separate outcomes by their subjective probabilities. Risk aversion is formally captured by a concave utility function. But, again, what adaptive purpose does concave utility serve? In fact, why even have a utility function at all? Robson (2001) answers this question in a setting where evolutionary fitness is determined by a simple two-armed bandit problem (e.g. a choice between which of two fields to forage in on a given day). In such a simple setting, it is optimal to choose the field with the higher probability of yielding the most berries. Therefore one might surmise that evolutionary forces would simply select for agents who always choose this field; essentially hard-wiring this action. But such a trait would clearly be sub-optimal in a world in which the fruitfulness of the fields can change. In this world, Robson shows that a fitness-maximizing agent should be endowed with a utility function over outcomes.[3] This is a functional argument for describing behavior consistent with utility maximization.

---

[2] Alternative approaches include expanding the functional question to include constraints on attention, as in the rational inattention literature (Woodford, 2020; Caplin & Dean, 2015), as well using implementational knowledge to inform the maximization algorithm (Fehr & Rangel, 2011; Webb, 2019; Landry & Webb, 2021). Hebert & Woodford (2021) and Frydman & Jin (2021) link some of these insights together.

[3] Moreover, the shape of the utility function is determined by the energy requirements required for procreation (or in the language of economics, the "baby production function" must be concave).

So how might utility-maximizing behavior be implemented at the algorithmic level? When agents do not know the probabilities in a bandit task *a priori* – and must therefore learn them – simple reinforcement learning (RL) algorithms offer computationally efficient solutions (Sutton & Barto, 1998). In a RL model, the agent samples a bandit and updates an associated value based on the outcome. This updating is implemented through a *reward prediction error*, reflecting the difference between the realized reward and the predicted value of the bandit. Eventually, as the agent repeatedly samples from each bandit, the associated value functions will converge to the true expected reward.[4] Remarkably, neuroscience research has shown that these reward prediction errors are encoded in the brain through the activity of dopamine neurons (Schultz, Dayan, & Montague, 1997; Glimcher, 2011). These studies, which form some of the foundational results in the field of neuroeconomics, thus describe the neural computations for implementing reinforcement learning — an algorithm that can generate behavior consistent with utility maximization.[5] In fact, binding together these restults at all three levels is considered one of the great successes of Marr's approach (Niv & Langdon, 2016).

In the following section, we will examine how these same neural signals are used to guide habitual behavior, in which an individual automatically repeats a previous action. While such thoughtless repetition of behavior may seem to diverge from principles of utility maximization, in stable reward environments it might be efficient to automate decision-making to continue choosing an alternative that has been optimal in the past, thus

---

[4] How well this algorithm performs depends on the properties of the environment (like the variance of its reward distribution, or in the case of foraging, the prevalence of extreme events like drought). But generally these algorithms perform well when agents are forced to explore with some small probability, ensuring that the agent does not initially reject a bandit simply because of a few poor outcomes. Remarkably, this algorthm also works well when the reward distribution is non-stationary; the exploration probability ensures that all bandits will be continuously sampled, implementing a tradeoff between exploiting learned information and exploring for more valuable actions. The simplicity of this algorithm, and its ability to guide behaviour in such tasks, has spawned the growing field of machine learning (Sutton & Barto, 1998).

[5] However, this system can also go awry. Anything that changes reward prediction error can change learned subjective values and causally affect choice. One amazing demonstration used optogenetics—using light to locally turn activity in small ensembles of neurons up and down. Steinberg et al. (2013) trained rats in a standard conditioning test pressing levers for rewards of sugar or water. By optically stimulating ventral tegmental area (VTA) dopaminergic neurons, they created artificial subjective RPE signals for low-reward outcomes. The rats behaved as if water was sugar. There is no way to explain this in conventional economic theory because it requires a mechanistic account of what brain area encodes RPE and how RPE's can be artificially created.

foregoing the costly cognitive processes of collecting utility-related information and consulting one's utility-ranking for the Nth time. However it would also be important to have a mechanism for turning off a faulty habit when conditions change. The goal of this chapter is to bind together these ideas and evidence on habitual choice at all three levels of Marr's analysis, while clarifying the broader goal of neuroeconomics as follows:

> **Neuroeconomics seeks to identify the algorithms of economic choice that achieve high-level functional goals, and are actually implemented by neural circuitry and other biological mechanisms (such as genes and hormones).**

II. Habit

A healthy person going about their day is constantly shifting between highly automatic, "habitual" behavior and interruptions for thoughtful deliberation. For example, as Sharon gets in the car for the Monday morning commute, she plunks her home-brewed coffee in the drink-holder, fastens her seatbelt, punches FM station setting 2, and backs out of the driveway. She has done this many times before, and she does so now with minimal attention and possibly without even remembering having done many of the steps. She is using habitual control.

On her way to work, suppose the radio station alerts her to terrible traffic on highway 134 in Burbank. Construction began last week and is expected to continue for two more weeks. If Sharon is deep in habit mode, she does not even process the radio announcement. She gets stuck in traffic and then encodes a negative "prediction error" (a gap between the driving outcome she experienced today and what she experienced in the past). After a few days of getting stuck in traffic – perhaps only one day, if she learns quickly – she will be jolted out of her habit and contemplate potential alternate routes for her commute.

But what if Sharon had only made this commute for a few days, and had not yet formed a habit? In this case, Sharon would process the initial traffic alert using a cognitive system that quickly sprouts decision trees representing different route options. The trees include a "staying on Highway 134" branch, which Sharon expects will lead to bad outcomes (such as "frustration", and "late for work") as a result of terrible traffic. Sharon's mind rapidly starts thinking through alternate branches and chooses a new route based on the expected quality of outcomes associated with these newly-imagined branches. She'll exit highway 134, take a Cahuenga detour, and stop at a favorite coffee shop. She can call into her 9 am meeting from there.

Our neuroeconomic theory of habit posits that behavior is the output of one two neurally-distinct systems for choice.[6] In *habit* mode, people who face a familiar choice situation make the same choice they did the last time if that choice was reliably rewarding. The habit mode is fast, implicit (i.e., often not conscious), and can make optimal choices in a sufficiently stationary world. By contrast, in a *model-based* mode people use a subjective representation (or model) of how their actions in different states correspond to likely outcomes, and the value of those outcomes.[7] This mode is slow, is more likely to be conscious (or "explicit" in psychological language, i.e., people can describe their behavior and associated thought processes), and more computationally costly (Daw et al., 2011; Daw & O'Doherty, 2014).

The concept of habit has, of course, been discussed for a long time in many disciplines, including economics, sociological anthropology ("rituals"), and individual and organizational psychology ("routines").[8] The Nobel laureate economist Gary Becker (Becker, 1996, p.9) wrote "I believe the main reason habitual behavior permeates most aspects of life is that habits have an advantage in the biological evolution of human traits." Here, we do not formally address the normative question of what a habit-based system

[6] "Dual system" neuroeconomic models have been proposed before. The closest in spirit to ours is Benhabib & Bisin's (2005) model of automatic and controlled interaction in savings decisions (and mention predictions about neural activity and response time). However, there is no dynamic concept akin to habit in their approach.

[7] In cognitive neuroscience, the "model-based" terminology is meant to distinguish this decision mode from simple reinforcement learning of the value of a state or action through direct reinforcement (which is perhaps confusingly called "model-free learning"). Here, the model-based mode can be understood to include models like expected utility maximization, where the association between states and outcomes are learned via Bayes Rule.

[8] See sociological discussions about norms and organizational "routines" (Nelson & Winter, 1982; Biggart & Beamish, 2003; Hodgson & Knudsen, 2004). Routines are company-level habits.

should look like.[9] Rather, our goal is to describe the behavioral predictions that result from the existence of these two decision-making modes (each of which have their own supporting empirical evidence in the human brain) and transitions between them. To this end, we make predictions about how consumers respond to incentives and disincentives as captured by a basic economic measure – price elasticity of demand.[10]

a. Neuroscientific background

The distinction between habit and model-based decision-making modes, as we refer to them, first gained precision in animal learning, then quickly became evident in a variety of neuroscientific studies.[11] This section will first sketch precursors to habit in biology, then present our definition of habit.

Automatic behaviors are ubiquitous in nature. Most animals have reflexes that are highly automated and do not require conscious thought. Lorenz (1935) showed that newborn ducklings will immediately bond with and follow around their mother as soon as they can walk— and if the mother is absent they will follow around a human (a phenomenon called "imprinting"). Animals exhibit fixed action patterns when confronted with an "innate releasing mechanism" in response to an evolutionarily-important stimulus (e.g., mating dances by birds; Tinbergen, 1951). Later scientific study, in simple domains such as rodents running through mazes for food, showed how chained sequences of actions become rapid and automatic. Neural activity is observably reorganized as sequences are learned (Graybiel, 1998; Smith & Graybiel, 2016; Smith & Graybiel, 2014) .

An important difference between reflexes and habits is that reflexes are typically difficult to control or override. Regular breathing is a reflex; babies are born doing it. The magician David Blaine trained himself to hold his breath for several minutes, and free

---

[9] However our conceptualization of habit is based on its potential efficiencies compared to a model-based choice (at least in some stable environments). For a model of habits as optimal responses to information processing frictions, see Matyskova, Rogers, Steiner & Sun, 2020.
[10] See Camerer et al (2021) and Xin et al (2021) for ongoing applications to consumer choice and social media use.
[11] Psychological ideas and measures of habits proceeded in a parallel track but have not been closely linked to biological mechanisms (Wood & Rünger, 2016).

divers (including some individuals from cultures with long traditions of underwater fishing) can too. But that type of override only occurs with long periods of training. Habits, in contrast, can be inhibited more easily.

The first step in our approach is forming a precise definition of habit, with empirical content. Adams & Dickinson (1981) defined a habit as a behavior that is so "overlearned" that it is repeated even when rewards are devalued. For example, Adams & Dickinson (1981) conducted a seminal study in which rats were trained to press a lever for sucrose pellets in ten training sessions, over 500 rewarding trials. Lever-pressing becomes rapid and automatic. After this training, the sucrose is paired with lithium chloride and the rats are allowed to eat as much sucrose as they want. This pairing makes the rats mildly ill and creates a taste-aversion so that they dislike the sucrose pellets. The sucrose has been "devalued." (Other methods of devaluation are to allow rats and humans to eat or drink to satiety, at which point the marginal value of additional reward is thought to be zero.)

After devaluation, the rats are allowed to lever-press again (with no reward presented). Rats press about 4.4 times per minute after sucrose was devalued, similar to 3.7 times per minute for control rats (for which there was no devaluation). Five hundred trials of training apparently built up a habit to press the lever, *even when pressing does not achieve a desirable goal*. In a group with less training (100 trials) rats in the devalued condition *did* press less often, only 1.7 times per minute, than control rats. Numerous replications of this classic result can be found in neuroscience and psychology, particularly for animals (Dickinson, 1985; Dezfouli & Balleine, 2012). However, the effect of over-training is more nuanced in humans (perhaps related to stress), even when a majority of subjects exhibit insensitivity to reward devaluations (i.e. habitization) (de Wit et al., 2018; Pool et al., 2021).

In microeconomic terms, the rats are persistently "demanding" food at the same rate, even after its quality has changed. This is naturally described as a zero quality-elasticity (an elasticity is the percentage change in demand relative to the percentage change in another variable, such as price or income). While price does not change in the animal experiments, it is a small leap to expect that a habit characterized by zero qualityinelasticity will also exhibit zero price-elasticity.[12] From a neuroeconomic perspective, a habit can therefore be *defined* by a repetitive choice pattern with perfectly *in*elastic demand. This neuroeconomic definition links the psychological concept of habit from animal learning with a behavioral measure from economics that specifies how a habit may be identified in economic data.

A related stream of research establishes a basic distinction between habit and "goal-directed" control – the latter idea mapping to our notion of model-based decision-making – and explored their neural mechanisms (Dickinson, 1985; Dickinson & Balleine, 1993; Yin & Knowlton, 2006). Some studies even seem to pinpoint the brain regions used to execute habits. If a habit is short-circuited, animals should cease to respond when reinforcers are devalued (i.e., they *will not* respond for worthless rewards, as in the experiments just described). Indeed, causally extinguishing activity in the dorsolateral striatum with neurotoxins erases habit in rats (Yin, Knowlton, & Balleine, 2004; Kyle Stephen Smith & Graybiel, 2014, Fig.1). The homologous area of the dorsolateral striatum in humans is also activated by fMRI when people execute habits (Tricomi, Balleine, & O'Doherty, 2009; Wesley et al., 2014).[13]

As Gary Becker's quote anticipated, a *mixture* of habit and model-based choice is an ideal way for healthy, busy, humans to behave. In developed economies, people make many choices every day, and only have a scarce amount of cognitive control (or "executive function"). Habits allow the brain to operate on "autopilot" and conserve mental resources (see Duhigg, 2012, for a good popular account). While habitual choices can be optimal if the subjective values of available goods are not changing too much, habits can overlook

---

[12] Note that from the functional point of view, it is perhaps no surprise that habit creates temporary insensitivity to devaluation. The cognitive economic advantage of habit is to reduce expenditure of scarce mental resources on regular, reliably-rewarding behaviors. The strength and automaticity of habit will naturally tend to lead to persistence of a habitual behavior when it is no longer rewarding.

[13] A high level of activity in this area during choices for drugs (compared to activity in control subjects), in concert with reduced dorsolateral pre-frontal activity, is linked to addictive behavior in humans. As Wesley et al. (2014) describe: "the evolutionarily young executive system [pre-frontal] … works in concert with the older limbic system [striatum] for optimal decision-making. When these systems are functioning sub-optimally decision-making becomes impaired." Note, however, that we are focusing on a model of healthy individuals, making relatively mundane, regular choices. See Section IId for further discussion of addiction.

opportunities to switch to choices that have become more rewarding.[14]

The key feature of a well-functioning habit is a mechanism for turning habits on and off. An approach from theoretical neuroscience proposes that the brain keeps track of the reliability of rewards from the individual's choices and uses these reliability measures to activate and deactivate a habit (so that both computations are being made in parallel; Daw, Niv, & Dayan, 2005; Keramati, Dezfouli, & Piray, 2011). Ideally, this arbitration should break habits when they become mistakes, and should hand control back to the habit system when choice values stabilize (i.e., reward reliability is high). Lee, Shimojo, & O'Doherty (2014) provide the first fMRI evidence of this process. They identified activity in lateral prefrontal cortex and frontal pole encoding reliability signals for both habitual and non-habitual choices. Furthermore, they find that the persistence of a habit is best explained by a reward prediction reliability measure formed from the absolute values of accumulated reward prediction errors. Notably, this suggests that a habit may be turned off by both good and bad surprises about the value of a habitually chosen option.

b. The neural autopilot model

Here we present a simple theoretical model that formalizes our neuroeconomic concept of a habit (and nests a standard utility-maximization model as a special case). A key feature of our model is how control over behavior is transferred from one decision-making mode to the other, what computations are made in each mode, and the choice predictions that result. We call the theory "neural autopilot" because it resembles a control system – such as flying an airplane – that automates some behaviors, freeing up the scarce cognitive resources of pilots to monitor and control non-automated activity of the plane.

[14] This same principle leads to what are called "action slips," in which a habit is executed without minimal cognitive resources, such as a tired professor putting ground coffee directly into a glass, rather than a coffee machine. Action slips also can take place in unfamiliar environments in which trigger cues or design has changed, even slightly. For example, author Camerer's household now has two SUVs. One has an old-fashioned physical key ignition system and the newer one has an electronic key fob so pressing a button starts the car. He often mixes up the two, pressing the keyholder in the older car to start it.

Such well-functioning systems have sensors that alert the pilot when an automated system might need to be disabled and overridden by human control.

Formally, a consumer faces a repeated choice between two options, A and B, at each period $t = 0,1,...$. We let c_t denote the consumer's choice (either A or B) at t and $u_t(x)$ denote the net subjective value or utility from choosing x at t. Following suggestions and evidence from neuroscience, the consumer forms *reward predictions* regarding the value of each option and updates these predictions based on values realized from the consumer's choices. Specifically, the reward prediction for x at time t, denoted $r_t(x)$, is modeled as a time-weighted average of past utilities from choosing x according to:

$$r_t(x) = \begin{cases} (1-\lambda_r)r_{t-1}(x) + \lambda_r u_{t-1}(x), & c_{t-1} = x, \\ r_{t-1}(x), & c_{t-1} \neq x, \end{cases}$$

where λ_r is a constant learning rate (between 0 and 1) that reflects the degree to which reward predictions adjust to new information; if $\lambda_r = 1$, reward predictions fully update to reflect the value realized from choosing the option in the previous period, but if λ_r is low, reward predictions update more gradually. Of note, the updating rule for $r_t(x)$ after x is chosen (i.e. when $c_{t-1} = x$) can be re-written as $r_t(x) = r_{t-1}(x) + \lambda_r(u_{t-1}(x) - r_{t-1}(x))$, in which case the new reward prediction is expressed as the previous reward prediction plus a term that represents a *reward prediction error* (RPE), i.e. $u_{t-1}(x) - r_{t-1}(x)$, scaled by the learning rate λ_r. If x was not chosen ($c_{t-1} \neq x$), there is no new information about its value, so the associated reward prediction does not change.

Besides forming predictions regarding the value of each option, we assume that the consumer also tracks the *reliability* of these predictions. We let $d_t(x)$ denote the "doubt stock" for option x at time t. A low value of $d_t(x)$ means the consumer implicitly "trusts" (i.e. has little doubt in) the current prediction $r_t(x)$. Doubt is low if past reward predictions have been good predictors of past utilities; doubt is high if past reward predictions were far off. Formally, and motivated by previously discussed findings from neuroscience, the doubt stock updates based on the *absolute values* of RPEs, $|u_t(x) - r_t(x)|$, as follows:

$$d_t(x) = \begin{cases} (1-\lambda_d)d_{t-1}(x) + |u_t(x) - r_t(x)|, & c_{t-1} = x, \\ (1-\lambda_d)d_{t-1}(x) + \alpha & c_{t-1} \neq x, \end{cases}$$

where λ_d is a constant learning rate between 0 and 1 that is similar to λ_r except it applies to the evolution of the doubt stock as opposed to the reward predictions. For simplicity, we assume $\lambda_r = \lambda_d = \lambda$ going forward, though empirically λ_d and λ_r could be different.

Our rule for updating the doubt stock indicates that if x was chosen at $t-1$, the doubt stock at t will be updated based on the *absolute* RPE, $|u_t(x) - r_t(x)|$. If $u_t(x)$ does not change as x is repeatedly chosen, then the absolute RPE will approach zero, causing the doubt stock to approach zero, but if $u(x)$ suddenly changes by a large amount, a large RPE will cause a jump in the doubt stock. If x is *not* chosen at $t-1$, however, its updated doubt stock is $d_t(x) = (1-\lambda)d_{t-1}(x) + \alpha$; here, the addition of the constant $\alpha > 0$ reflects a rise in doubt because the reward prediction for x was not actually "tested" against its true utility. Going forward, we normalize $\alpha = 1$. If x is repeatedly not chosen, its doubt stock will converge to λ^{-1}. We therefore set $d_0(x) = \lambda^{-1}$ for $x = A, B$, which ensures the initial doubt for a given option will be constant as long as that option is never chosen.

While the impact of the RPE on doubt could, in principle, be stronger or weaker depending on whether the RPE is positive or negative, in our model positive and negative RPEs are weighted equally. This assumption may in part be understood as a simplification for analytical convenience. With that said, it also operationalizes previously discussed evidence from neuroscience that absolute values of RPEs may be used to form reliability estimates that govern the use of habit.

We now define our two decision-making modes. In *habit-based* decision-making, the consumer simply repeats her previous choice:

$$c_t = c_{t-1}.$$

By contrast, in *model-based* decision-making the consumer considers both options and chooses the option that maximizes present utility:

$$c_t = \begin{cases} A, & u_t(A) > u_t(B), \\ B, & u_t(A) < u_t(B). \end{cases}$$

Here, we assume the consumer considers both options and uses all utility-relevant information in a given period to inform their choice; in this way, our formulation of model-based decision-making corresponds to a standard microeconomic approach.

At this point, one may wonder why a consumer would bother forming a reward prediction r_t if the consumer can just observe u_t. First, reward predictions are still useful in other environments where it would *not* be possible to know u_t in advance of a choice. Recall from our earlier discussion, reward predictions are used in reinforcement learning models that provide a simple algorithmic solution to such problems (e.g. multi-armed bandits) and that their neurobiological markers form the basic reward learning circuitry in the brains of vertebrates. Even in environments where it is possible to observe utility relevant information, it has been demonstrated that reward predictions are still computed in the brain even if they are ultimately ignored (Lee, Shimojo, O'Doherty, 2014). Moreover, our model may be understood to include a fixed cognitive cost that is paid in each period that the consumer uses model-based decision-making. The existence of a habit system is premised on the inefficiency of using model-based decision-making in all choice environments. But the decision-maker then needs a mechanism to arbitrate between modes. As we will soon see, reward predictions are a key component of this mechanism.

To close the model, habit-based decision-making is used if and only if the doubt stock associated with the previously-chosen option is below some threshold θ with $0 < \theta < \lambda^{-1}$. This leads to the following choice rule:

$$c_t = \begin{cases} c_{t-1}, & d_t(c_{t-1}) < \theta, \\ \arg\max_{x \in \{A,B\}} u_t(x), & d_t(c_{t-1}) \geq \theta. \end{cases}$$

If $\theta = 0$, the consumer always uses model-based decision-making. In this way, our model nests a standard utility-maximization model as a special case. Even with $\theta > 0$, the consumer necessarily uses model-based decision-making at $t = 0$, choosing the option among A and B with higher utility – here, the use of model-based decision-making is assured since $d_0(A) = d_0(B) = \lambda^{-1} > \theta$.

If $u_0(A) > u_0(B)$, in which case the consumer's first choice will be $c_0 = A$, and utilities do not change, the consumer will continue to choose A while the doubt stock $d_t(A)$ decreases over time. Once $d_t(A)$ falls below θ, habit mode takes over and the consumer continues to choose good A *without even considering good B* (hence the concept of autopilot). Meanwhile, habit mode will control choice until it is interrupted by a series of sufficiently large RPEs, caused by changes in A's utility, pushing $d_t(A)$ above θ. Recalling that the doubt stock is driven by *absolute* RPEs, we can see that both "good" ($u_t(A) > r_t(A)$) and "bad" ($u_t(A) < r_t(A)$) surprises can jolt the consumer out of habit mode, prompting model-based consideration of both options. Thus, even if good A's price or quality improves, this could lead the consumer to abandon the habit and choose B.[15]

Stepping back, we can now see how the model clarifies the reduced-form economizing value of habit. A computationally cheap process that determines the decision-making mode first perceives an environmental state (including cues which may be associated in memory with previous choices in that state[16]), recalls the previous choice c_{t-1}, and checks if the doubt stock $d_t(c_{t-1})$ associated with this previous choice is less than θ. If the answer is "Yes" the previous choice is executed out of "habit." Here, a habituated person *does not even need to gather information or recall the utility of what they are buying*. In terms of neural activity, we would expect to find neural signals associated with memory of choices and with value uncertainty (i.e. "doubt") in habit mode. There are likely to be weaker signals associated with reward prediction and with subjective value.

c. **Habit and price elasticities**

Next we show two interesting implications of our habit model for price elasticities (defined as the percentage change in quantity demanded divided by the percentage change in price). Price elasticities are represented as unitless numbers, allowing them to be compared across different types of goods. Elasticities are workhorse measures in economics, used to

[15] Here is a motivating example of how a positive RPE might shift people out of habit mode. Suppose a family goes to the same Chinese restaurant every Sunday and orders their same, favorite dishes. One Sunday two of the dishes are unusually good. Intuitively, even though the dishes taste better, positive prediction error increases doubt stock (perhaps triggers in an inference that there is a new chef) and could lead to deliberate evaluation of other dishes they did not like in the past. In other words, the positive surprise provides information that the decision-making environment might have changed (i.e. the distribution of rewards in an environment are not entirely independent). Our conjecture is that the human decision-making system may have evolved under such conditions (e.g. societal and climatic).

[16] Situational cues have been shown to be important in biological addictions, such as opiods. The theory is that homeostatic processes designed to maintain stability in physical states respond to Pavlovian cues which predict immediate drug use in a "feed-forward" process. Siegel et al (1982) showed that rodents exposed to a series of increasingly-large heroin doses were twice as likely to do from a large novel dose when the dose was administered in an unfamiliar environment. Laibson (2001) explores some formal implications of a model with cue-dependence. Wellsjo (2021) shows that unfamiliar situational cues (e.g., working in a different part of a hospital) affect habitual use of handwashing mchines.

define monopolization (low elasticity) and competition (high product-specific elasticity), to define substitute and complementary goods (cross-price elasticities that are positive and negative, respectively), and to quantify the impacts of policy changes such as increases in gasoline taxes.

An ongoing challenge for neuroeconomics is whether understanding basic neural computational processes – and constructing models about those processes (at Marr's algorithmic level) – can say something novel about core economic questions. We developed the autopilot model specifically to say something about a very basic economic construct: price elasticity. As addressed earlier, there is a simple link between neuroscience and elasticity, through our concept of habit. As we recall, Adams & Dickinson (1981) conceive habit as insensitivity to devaluation of outcomes, which naturally maps to demand characterized by zero price (or quality) elasticity.

Next, we explore our model's implications for price elasticities in a stylized consumer market, focusing on the anomalous responses to price changes among consumers who have already formed habits. We assume the market is comprised of a unit mass of consumers with symmetrically heterogeneous preferences in that, with equal (initial) prices for A and B, i.e. $p(A) = p(B) = \bar{p}$, the utility difference $u(A) - u(B)$ is presumed to be uniformly distributed between -1 and $+1$ across all consumers.

Next, we consider the effect of a change in the prices of A and B by Δ^A and Δ^B, respectively (so that $\Delta^x > 0$ means x has become more expensive while $\Delta^x < 0$ means x has become less expensive). Letting $\mathrm{Q}^A(\Delta^A, \Delta^B)$ represent the quantity demanded for A in the period when these price changes first occur, the *own-price elasticity* for A quantifies the effect of the price change Δ^A on the demand for A and is given by:

$$\eta^A(\Delta^A|\Delta^B) = \frac{[\mathrm{Q}^A(\Delta^A,\Delta^B) - \mathrm{Q}^A(0,\Delta^B)]/\mathrm{Q}^A(0,\Delta^B)}{\Delta^A/\bar{p}}.$$

The *cross-price elasticity* quantifies the effect of B's price change, Δ^B, on the demand for A:

$$\eta_C^A(\Delta^B|\Delta^A) = \frac{[\mathrm{Q}^A(\Delta^A,\Delta^B) - \mathrm{Q}^A(\Delta^A,0)]/\mathrm{Q}^A(\Delta^A,0)}{\Delta^B/\bar{p}}.$$

From here on, the arguments (Δ^A and Δ^B) in $\eta^A = \eta^A(\Delta^A|\Delta^B)$ and $\eta_C^A = \eta_C^A(\Delta^B|\Delta^A)$ will often be suppressed for ease of exposition.

There are four qualitatively distinct ways in which demand in a market with habituated consumers can depend on price changes, each of which represents a different-colored region in the following graph:

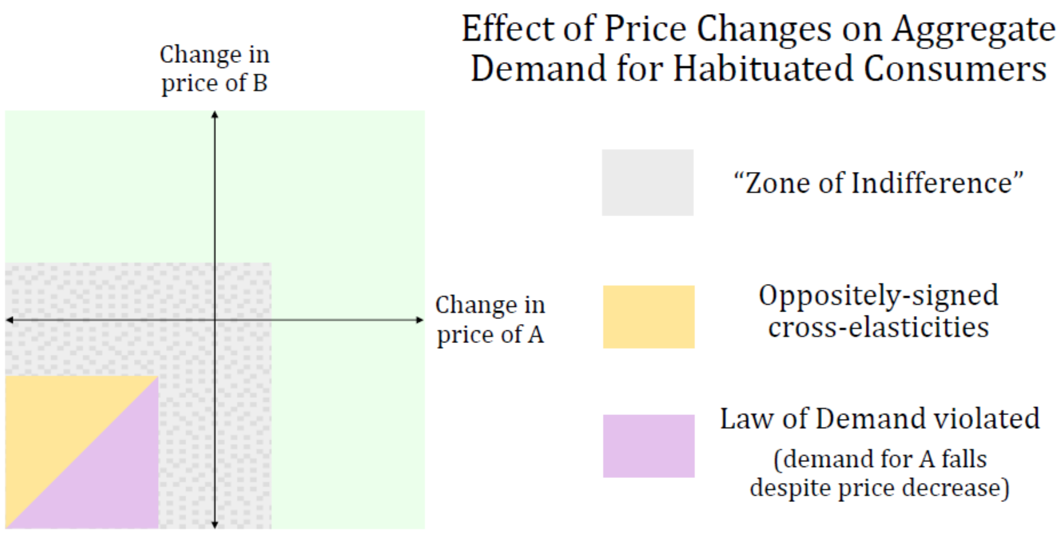

Figure 1: Changes in demand for A as a function of price changes of A and B in a market of consumers with pre-established habits (half for A, half for B).

The green area in Figure 1 represents a "normal" region in which price changes induce the same "standard" effects on demand for each good as for consumers in the model-based decision-making mode. Here, consumers' aggregate behavior satisfies the "Law of Demand" in that demand for A falls as its price increases ($\Delta^A > 0$) and rises as its price decreases ($\Delta^A < 0$), as captured by a negative own-price elasticity, $\eta^A < 0$. Furthermore, demand for A rises with an increase in the price of its substitute B ($\Delta^B > 0$) and falls with a decrease in the price of B ($\Delta^B < 0$), as captured by a positive cross-price elasticity, $\eta_C^A > 0$ (the converse also holds: $\eta_C^B > 0$).

The three remaining regions highlight ways in which habituated consumers may exhibit nonstandard responses to price changes, as predicted by our model.

Zone of Price Indifference (*Gray*). Within the gray region, the price changes Δ^A and Δ^B have no effect on demand; that is, *price-elasticities are zero*: $\eta^A = \eta_C^A = 0$. This so-called “zone of price indifference” (ZPI) exists because price changes within the zone are either not large enough (in magnitude) to jolt consumers out of habit mode, or because a price change is sufficiently favorable to preclude switching to the other option even when both options are considered. For example, a huge price decrease for A can prompt consumers habituated to A to consider both options, but since A is even more appealing in relation to B at their new prices in this region, those who have been consistently choosing A will continue to do so – except now based on a higher valuation of A rather than a habit.

In line with this prediction, many marketing studies have documented near-zero price elasticities in a measured ZPI surrounding a reference price.[17] Estimates of the size of the ZPI are about 10% of the current price (Han, Gupta, & Lehmann, 2001). They cluster all consumers into two segments based on several observables. The two segments appear to correspond to more price-sensitive and more habitual shoppers.[18]

Law of Demand Violated (*purple*). According to our model, the Law of Demand can be violated under the right conditions among consumers with pre-established habits. Namely, within the purple region of Figure 1, $\eta^A > 0$, as a decrease in the price of A paradoxically leads to a decrease in demand. In this region, a decrease in the price of A is sufficiently large to induce a shift to model-based decision-making, but insufficiently large to compensate for the even larger price decrease for B, leading some consumers to switch from A to B. If A had maintained its original price, however, these consumers would not have noticed B’s reduced price and would have therefore continued to choose A (even

[17] The ZPI is also called “latitude of price acceptance” in marketing. For evidence of a ZPI, see Kalwani and Yim (1992), Gupta and Cooper, Kalyanaram and Little (1994), Han et al. (2001), and Terui and Dahana (2006), among others.
[18] Their algorithm compares fits for two to five clusters and picks out two as the best-fitting (meaning additional clusters did not change overall fit much). Other studies have found two or three segments (Bucklin & Gupta, 1992). This is consistent with a two choice-mode simplification. While not addressed by our analysis of a one-time price change, there is also mixed evidence that indifference thresholds increase when prices are more volatile, consistent with “hysteresis” stemming from the option value of waiting for a better deal when prices are variable (Richards, Gómez, & Printezis, 2015).

though B now represents a better deal). Notice that this violation of the Law of Demand arises even though the options in our model would not be considered "Giffen goods."

Such incorrectly-signed own-price elasticity estimates – indicating violations of the Law of Demand – are indeed common in the empirical literature (e.g. as discussed by Montgomery & Rossi, 1999). However, such estimates are generally not accepted at face value because they are thought to lack a viable theoretical foundation (also see Cramer, 1973; Song & Chintagunta, 2006). In this light, our model offers a potential theoretical rationale for this phenomenon.

Oppositely-Signed Cross-Price Elasticities (*orange*). Standard economic treatments suggest that the relationship between the price of B and demand for A should be qualitatively the same as the relationship between the price of A and demand for B. Namely, when A and B are substitutes (as in our model), a price change for one good should shift demand for the other good in the same direction. For instance, if the price of Coke goes up, demand for Coke goes down (the own-price elasticity is negative) while demand for Pepsi goes up (the cross-price elasticity of the substitute good is *positive*).

While in principle cross-price elasticities η_C^A and η_C^B should both be positive for substitutes, within the orange region of Figure 1 the cross-price elasticities have opposite signs, $\eta_C^A < 0 < \eta_C^B$, as the price decreases $\Delta^A < 0$ and $\Delta^B < 0$ both shift demand towards A and away from B. The reason B's price cut shifts demand towards A in this case follows from the fact that the Law of Demand for B is violated within this region (not to be confused with the purple region, where the law is violated for A), leading to the peculiar effect on demand for A. As noted in a review by Bonfrer, Berndt, & Silk (2006), empirical studies frequently find that estimated cross-price elasticities in pairs of goods have opposite signs. There is no consensus about how asymmetries should be interpreted, or whether they reflect measurement error. The autopilot theory gives one possible explanation.[19]

[19] In a meta-analysis of 15 studies on grocery purchases, Sethuraman, Srinivasan, & Kim (1999) report that about 10% of 1,060 cross-price elasticity estimates are incorrectly signed. The authors suggest these incorrectly-signed elasticities were probably due to measurement error. We believe this explanation is probably correct in many cases, but it is also possible that anomalous findings such as

Asymmetric Responses to Gains and Losses. The model can also generate a gain-loss asymmetry in that demand is more responsive to increases in the price of good A than it is to equal-sized price decreases. This prediction is most apparent when A's price change is sufficiently large to escape the ZPI for a price increase $\Delta^A > 0$ but $|\Delta^B|$ is small enough to keep consumers in the ZPI for a price decrease $-\Delta^A < 0$ of equal magnitude. The intuition is that both an increase or a decrease in A's price create RPEs that are equal in absolute value. Both changes will push the doubt stock $d_t(A)$ above the threshold θ, triggering model-based decision-making – but only among consumers who were habituated to A. While the price increase will shift demand away from A (among consumers previously habituated to A), a price decrease would not draw a similar share from consumers habituated to B if $|\Delta^B|$ was too small to induce these consumers to consider A at its new price.

Such asymmetry in consumers' responses to price changes was documented in Putler's (1992) empirical analysis of the demand for eggs. He estimated elasticities for price increases and decreases to be -.78 and -.31, respectively, while many other studies have confirmed this asymmetry (e.g. Mayhew & Winer, 1992; Hardie, Johnson, & Fader, 1993; Kalyanaram & Winer, 1995; Han et al., 2001). The standard behavioral explanation is that loss aversion from paying a price higher than a reference price causes this asymmetry. This explanation is certainly plausible (and is consistent with other evidence of loss aversion in consumer behavior e.g. Ahrens, Pirschel, & Snower, 2017). However, our autopilot model can generate the same theoretical result without loss-aversion. Furthermore, our model makes a new empirical prediction: that asymmetric responses to price changes will only be seen for large enough price changes in one product while any price changes for substitute products are small. If both changes are large enough, there will be no gain-loss asymmetry as Putler and others found.

d. Comparison with economic approaches

these are too quickly dismissed because there is no theory about why signs would be wrong. We offer such a theory.

In economics, the standard way to model habits is by assuming that past and present consumption are *complements* in the sense that past consumption of a good strengthens a consumer's current preference for that good (e.g. Pollak, 1970; Ryder & Heal, 1973; Becker & Murphy, 1988; Crawford, 2010). Such "preference complementarity" is typically formalized through utilities that depend on past choices; in the present framework, this could involve allowing $u_t(A)$ and $u_t(B)$ to depend on c_{t-1} while assuming that the difference $u_t(A) - u_t(B)$ is larger (all else equal) if $c_{t-1} = A$ than if $c_{t-1} = B$. While this approach captures a basic notion of choice persistence in habits, it does not tell us much about habits at the mechanistic, algorithmic, or even functional levels.

The preference complementarity approach has famously been used by Becker and Murphy (1988) to model the consumption of addictive goods, such as cigarettes. While there are several interesting parallels, our neural autopilot model is not tailored for the analysis of biologically addictive behavior. Our consideration of a regularly occurring choice between two options is instead more suited to study habits in consumers' repeated brand choices, such as a choice between two brands of cereal during a weekly visit to a grocery store.

By contrast, the frequency with which an individual considers consuming an addictive good likely varies substantially depending on the individual's addiction status. A full-fledged cigarette addict may contemplate smoking dozens of times a day yet 25 years after quitting this ex-addict may not even think about smoking on a typical day. Furthermore, in the context of addiction the decision of interest is generally whether or not to consume a particular good, as opposed to a choice between two versions (or brands) of the same good. As demonstrated in Landry (2019), a model of addiction based on such "endogenous consideration" can account for many empirical regularities that are not addressed by models of addiction based on preference complementarity.

The preference complementarity approach is closely related to representations of "state-dependence" (e.g. Keane, 1997; Dubé, Hitsch, & Rossi, 2010) – in which preferences for a good are affected by the previous purchase history (summarized as a "state" variable)

– for modeling consumer "inertia" or "brand loyalty" in brand choice.[20] Indeed, consumer purchase data usually show strong patterns of brand loyalty. While consumers in our model will exhibit a form of brand loyalty driven by a lack of consideration for other options while in habit mode, brand loyalty is not built in as a primitive. That is, there is no direct change in subjective value or choice frequency based on purchase history.

A particularly interesting study is Bronnenberg, Dubé, & Gentzkow (2012). They studied US consumers' purchases after moving to a new state. They find substantial persistence in the brands consumers buy: In the first four years after moving, consumers respond to just 60% of the new supply conditions they face (implying a 40% habit 'strength'). The implied habit strength is weaker for people who moved when they were younger.[21] Consistent with this finding, brand loyalty is generally shown to be stronger among older consumers (Lambert-Pandraud & Laurent, 2010).

The role of the threshold θ in our approach is similar to threshold approaches used in other areas of economics, where there is behavioral stickiness or rigidity absent a sufficiently large change in market conditions (Slade, 1999; Caballero & Engel, 1999; Cecchetti, 1986). Rigidity in this context is usually ascribed to frictions or "menu costs." Our approach is also related to macroeconomic work on "inattentive consumers." For example, Reis (2006) derives how frequently consumers re-evaluate their consumption plans when planning is costly. Consumption during periods of inattention is similar to habit, but decisions to plan are carefully computed by optimization. In contrast, we have opted to make the habit mode as "dumb" and automatic as possible.

Of course, there is plenty of empirical evidence of choices that *seem* habitual, in the sense that people persist in a regular choice even when it does not accomplish their desired

[20] Inertia and brand loyalty are also often conceived in terms of psychological "switching costs" (e.g. Farrell & Klemper, 2007). Switching costs could also result from learning-by-doing; for example, driving an automatic transmission long enough makes it difficult to switch to a manual stick shift. In these cases the increase in utility of a habitual action would naturally be interpreted as a reduction in effort cost. In a somewhat related approach – that perhaps more closely resembles the idea of a "hedonic treadmill" – in macro-finance "habit formation" preferences take the difference between a household's aggregate current consumption level and a weighted average of previous consumption (Abel, 1990; Constantinides, 1990; Boldrin, Christiano, & Fisher, 2001; Campbell & Cochrane, 1999).

[21] There is similar persistence in Indian immigrant preferences for foods from their birth regions (Atkin, 2013) and in European immigrant preferences for redistribution based on norms in their birth country (Luttmer & Singhal, 2011).

goal, or when it leads to a regretted mistake. For example, studies of people driving in simulators for many hours (Charlton & Starkey, 2011) demonstrate that highly-practiced drivers become inattentive. And in actual driving, near-crashes and crashes are often associated with inattention (NHTSA, 2006).

Defining habit purely in terms of preference complementarity has certainly been useful for some specific purposes. In our view, however, there are three limits to this approach that our neural autopilot model is intended to extend. First, the preference complementarity approach gives no guidance about what variables cause habitual choice to end. Our approach emphasizes the general two-controller model of choice and substitutes persistence of *choice* (during habit) for persistence of consumption driven by utility. This shift also invites consideration of questions like, "given environmental stationarity and action costs, what is the ideal mixture of habit and model-based decision-making?" This question does not make sense if habit is driven by preferences.

Second, if current consumption utility is increasing in past consumption, then reward signals associated with consumption should get larger and larger as the same good is chosen repeatedly (assuming utility generates reward signals). We know of no such evidence that persistent consumption increases reward signals (outside of a learning process), despite many, many studies on the learning of reward, ranging from direct measures of neural firing in monkeys to human fMRI.[22] This idea of increased reward (similar to "sensitization" in psychophysics) is not too biologically plausible as a general principle.

Third, a combination of evidence and intuition suggest that the development of habit can be identified empirically by multiple behavioral and cognitive markers (see Table 1). This is a large empirical advantage of our approach. The preference complementarity approach does not say anything about these markers because they do not specify the algorithm or mechanism by which habits work.

Establishing habit through insensitivity to devaluation (as in the Adams-Dickinson rat experiments) has been challenging in humans (Pool et al., 2021). Being able to identify

[22] Admittedly the studies that have investigated the effects of persistent consumption have used rewards that are more basic (i.e. food or drink) since they are typically run in animals and undergraduate subjects.

habit with a variety of measures is one way to work around this challenge; note that many of these markers are also part of implicit (or unconscious) processing. While not conventional, our view is that studying these empirical markers can enhance economic analysis. It is a way to use more data to decide between theories of choice that economists routinely use for basic purposes and policy analysis. Doing so does makes economics more like a natural science, which we think is good. Additional empirical guidance should be welcomed.

Table 1: Empirical markers of habit vs. model-based choice

	Habit	**Model-based**
Own-Price Elasticities	positive, negative, or zero	negative (has "correct" sign)
Cross-Price Elasticities	positive, negative, or zero	positive (for substitutes)
Choice Speed	Fast	Slow
Attention	Low	High
Purchase memory	Inaccurate	Accurate
Explicit explanation?	No	Yes

The first two empirical markers, relating to *own*- and *cross-price elasticities*, have already been discussed.

Animal and human studies also show that habitual choices tend to be made quickly (***fast response times***). This is a key property because response times are easy to measure. If habits aren't faster they aren't worthwhile. In addition, we conjecture that ***less attention*** is used in habit mode. This is uncharted territory, for the most part, because the neuroscientific approach to habit came largely from animal learning in a largely behaviorist tradition. While attention has not been measured in most of the studies that show consumption persistence, indirect evidence indicates that people often do not shop around much. For example, in the De los Santos, Hortaçsu, & Wildenbeest (2012) study of internet book shopping, 90% of people look at only one site within a day. Even in a seven-day window, 76% look at only one site. Marketing studies beginning in the 1970s show singleretailer shopping and low information acquisition is common, even for cars and expensive appliances (Newman & Staelin, 1972).

We also expect that people will ***remember*** habitual choices less accurately (since inattention tends to limit memory creation). An example of this memory effect is checking mobile cell phones. In one study people remembered having checked their phones 37.2 times/day, but the actual measured rate was 84.7 (Andrews, Ellis, Shaw, & Piwek, 2015). An anecdotal example of this habit amnesia comes from author Camerer's partner. She often orders online deliveries from Amazon. She was shocked when Amazon sent a congratulatory note about how many purchases she had made over several years—1500. Remember that in our neural autopilot model, the consumer does not need to keep track of the purchase history.

Note that different behaviors we think of as habitual can vary a lot on these psychological dimensions. Going to the gym is likely to be a combination of an explicit deliberation, followed by some low-attention choices about what gym clothes to bring, an automatic route to a locker room to change clothes, and then to a favorite treadmill. Ideal data from physical sensing as well as real-time self-reports would probably reveal a sequence mixing habitual and non-habitual choice, even for a single trip to the gym.

Finally, in most of cognitive neuroscience habits are thought to be ***implicit, not consciously known*** (Seger & Spiering, 2011). We conjecture that if you asked someone *why* they made a particular habitual choice, they would find it hard to explain in depth. They might answer superficially that "I usually do it this way" and "it's been good in the past." In a sense, those answers – as vague as they are – are actually a good reflection of cognition during habitual choice. "I usually do it this way" refers to the recall of previous choice. "It's been good in the past" is a simple way of saying "the doubt stock was low." In contrast, a model-based choice of, say, a home or a dinner at a new restaurant would probably refer to attributes of the choice, prices, the qualities of alternatives, and so on.

The last two empirical markers of habit– poor memory and explanatory depth – may not seem to be important economic variables. But memory and explanation are often inputs to other kinds of economic decision making. For example, many household surveys rely on memories of purchases. The Consumer Expenditure Survey (CES) uses both quarterly surveys and weekly "diary recordings" to measure large and small expenditures.

The CES is economically important because it is an input to inflation estimates (CPI). While households are encouraged in their diaries to record everything they buy, it is conceivable that habitual purchases are underreported. Understanding more about the habit-memory link will place boundaries on how common these biases are and help limit them.

Explanations of reasoning for choices are important too, for educational, organizational, legal, and political purposes (Willis & Reyes, 2015). In parent-child, teacher-apprentice, and manager-trainee relationships, if the more experienced person is making habitual choices and cannot explain why, their ability to foster human capital development by teaching junior learners is limited. In court proceedings, explanations of why defendants took certain actions can have important legal consequences. Thus, while psychological markers of habit like decision speed, memory, and explanatory quality are certainly not the typical variables of economic interest in studying decision making, we hope to measure them and explore how they could be useful.

We conjecture that there are two other empirical indicators of habit (though we have not derived these formally): search costs, and habit interruptions.

Search (and switching) costs: In habitual choice, by our definition, consumers are not searching among a large set of choices which may be varying in price and quality.[23] As a result, during habitual search the foregone value of search will be high. Any empirical analysis which infers the costs of search from actual purchase patterns will therefore infer that search costs must be high (to reconcile the low amount of search with optimization).

Many studies have tested models of optimal search using a combination of prices, choices, and search data. In these models there is a hidden "search cost" and consumers are assumed to search optimally. A habituated consumer will not search at all; the model will therefore identify habit in the form of a high inferred search cost.

For example, De los Santos, Hortaçsu, & Wildenbeest (2012) estimate the cost of searching one extra online bookstore for a bestseller as $4.14, assuming shoppers must search to learn prices (p. 2978). Hong & Shum (2006) estimate that half of subjects don't

[23] Turning to search raises the question of how our approach is related to "satisficing" (see the clear experimental evidence from (Caplin, Dean, & Martin, 2011). Satisficing is truncation of a search process after reaching a utility goal. Habit is shifting from a model-based decision process when repeated choice rewards are stable enough. There may be some interesting formal familial relation between the two concepts that would be useful to establish.

bother to search online for expensive textbooks; they buy at the first price. The estimated search costs for those who do search are about $1.30-$19 for the first additional search. Moraga-González, Sándor, & Wildenbeest (2013) estimate a search cost of around $8 for computer chips. If an extra search is quick, this implies a high marginal search cost per hour. Using data on exact shopping cart locations in a store, Seiler & Pinna (2017) estimate a return from search of $2.10/minute (or $126/hour).[24] De los Santos, Hortaçsu, & Wildenbeest (2012) also estimate rather low own-price elasticities for specific bestselling books at Amazon (from –.11 to -.65), while controlling for consumer knowledge of shifting price distributions. The Amazon elasticities are lower than other bookstores, which is consistent with habit if consumers use Amazon more regularly and form a habit.

It is also typical in these empirical search papers to find bimodality: A large fraction of consumers don't search beyond the first price, while others search more intensively. This is consistent with the hypothesis that consumers' choices can be guided by either of two decision-making modes. Of course, it is also likely that there are individual, experiential, and lifecycle differences in search intensity, namely income and time constraints. But any such theory would have to explain bimodality (e.g. why does the distribution of search costs have two modes rather than say a lognormal distribution?).

Several other economic papers study "inertia" in consumer choice. These papers do not fit models of optimal search; instead, they usually compare the costs and value of current purchases with opportunity costs from switching. It is possible that inertia is due to habit much as we model it, though these studies do not estimate habit explicitly. Luco (2019) gets empirical leverage from comparing new and regular customers in Chilean pension plan enrollment. He estimates an inertia cost of $37.50/month. Other estimated inertia costs seem to be high: $280 (online grocery delivery; Goettler & Clay, 2011), $2500 (health insurance; Handel, 2013) and $4000 (Medicare; Nosal, 2012). Hortaçsu, Madanizadeh, & Puller (2015) estimate it would take Texas households only 15 minutes to

[24] Intriguingly, they also find heterogeneity in imputed search costs for different locations in the store. This implies that product location is an important aspect to consider in habit formation. Consistent with some animal learning studies, even human shopping habits might be specifically constrained in time and space, or to particular motor actions.

save $100/year on electricity by switching. They also find that having an increase in last month's electricity bill increases search for a new supplier, which is consistent with the idea that large reward prediction errors jar people out of habit.

Habit interruption: An implication of our habit concept is that exogenous interruption of habits, typically through product unavailability, can trigger a model-based search. This possibility is consistent with evidence that "habit disruption" creates an opportunity for behavior change (Carden and Wood, 2018).

An intriguing analysis of the result of interruption is based on a two-day strike by workers in the London subway (Tube) system in February 2014 (Larcom, Rauch, & Willems, 2017). This is a good natural experiment because the strike only closed part of the Tube routes. Commuters affected by strikes could then be compared with similar commuters who were not affected. Oyster swipe cards. which store value and are used by many commuters to make cashless transactions, make it easy to record when and where commuters got on and off trains. Furthermore, the London Tube maps are not drawn to scale (they were created to look like compact circuit diagrams). As a result, a commuter who must get out at a new Tube stop and walk the rest of the way to work cannot judge precisely from the map how far they need to walk. They must learn, to some extent, from trial-and-error. In terms of our neuroeconomic model, the "inaccuracy" of the map will generate larger RPEs in learning for many commuters who break their habits and try out alternate routes.

During the strike, about 63% of commuters could not take their regular route. Consider habitual commuters, who had taken the same route for all 10 weekdays before the strike. After the strike, 5.4% more of them switched to a new route (compared to commuters who hadn't been affected by the strike and did not switch).[25] The switch saved an average of 40 seconds of time. Extrapolated over several years of commuting, (Larcom et al., 2017) estimated that the subjective search cost required to rationalize the fact that these commuters had stuck with a regular commute, but immediately switched because of the strike, was **£**380 ($636 in 2014).

---

[25] Note that creation of a new habit after an interruption is evidence against the common economic assumption that habits are a reflection of an increase in utility from making the same choice persistently.

The strike study raises an interesting question about how technology, regulation, and governmental action could enhance welfare with habituated consumers. The authors of the Tube study actually suggest that because commuters appear to search too little, the interruption due to strike was welfare-enhancing for commuters as a whole. This is a dramatic claim because it implies that in general, government actions which specifically make products or services unavailable for a short time, forcing search, can help people make better decisions by forcing exploration.

The welfare economics of habits is complicated and we have nothing too deep or special to say about it. There might be more efficient regulatory ways to improve consumer choice in the face of habit than *ad hoc* practices like occasional service-worker strikes. Regulation could force or subsidize exploration of new products, although competition in private markets may do so adequately without government intervention. Some technological solutions which enable consumers to search quickly and in a personalized way could also interrupt habitual choice and inform those consumers about better products. However, others might reinforce habitual behaviour, like social media and online shopping.[26] We do not have anything evidence-based to say about this process.

What happens in a market when firms know that consumers are habitual is an interesting question too. Because habitual choice is price-inelastic, by our definition, profit-maximizing firms have an incentive to retain habituated consumers or to help non-habituated consumers establish new habits. It is also true that if a firm continually raises prices or reduces quality by small enough increments, a habituated consumer may continue to buy the same product indefinitely. The model is not explicit about what other ingredients need to be added to counteract this predicted behavior. One plausible possibility is that habituated consumers pay at least a little attention at the point of purchase, enough to be sensitive to highly salient promotions for alternative products.

[26] Habitual choices aided by technological innovation may create a sort of dependence which have certain advantages (like saving time), but people, or even society as a whole, might also "unlearn" to do simple things themselves (e.g. dishwashers, microwaves, food delivery apps). This may contribute to an unsustainable lifestyle (at least when we consider the consumption side of this technical progress). That is, there may be positive or negative externalities of habitual behaviour.

Another consequence of habit is that new and improved products may have a hard time breaking into a market dominated by habitual consumers. Profit-maximizing firms with a good understanding of human nature should optimize against this constraint. However, if there are societal spillovers that cannot be monetized by a specific firm – such as adoption of cheap solar energy – there may be a role for government regulation, and recommendations will require good theories of consumer choice. In general, as in other areas of policy design, it will surely be useful to have a better understanding of the machinery which policies strive to influence – namely, the consumer or citizen brain – to figure out how firms and governments can improve welfare of habit-prone consumers, without removing the mental benefits of efficient habit formation.

Finally, the existence of a habitual consumer changes our view on consumer sovereignty and the "liberal" market conception, where the consumer reigns according to their well-thought-through preferences. If many of the consumers' choices are habitual, and firms know this and can influence new habits, that is, create a form of dependence between the consumer and the goods or services they buy, then the picture of power distribution and sovereignty changes dramatically.

Conclusion

We described a simple "neural autopilot" model of habit inspired by evidence from neuroscience and animal learning. The goal of models such as these is to have a sensible functional interpretation (what are habits *for*?), a tractable algorithmic specification which can be used to fit data and potentially prove theorems, and an underlying mechanistic implementation. In the model, habits are executed by recalling the previous choice (made in a particular environment or state) and a single measure of "doubt," and repeating the previous choice of doubt is low. If doubt is high, the decision maker reverts to model-based decision-making and considers both available options.

The model is used to make predictions regarding the demand-side responses to changes in price (or quality), as captured by price elasticities. A key insight from neuroscience is that habits entail insensitivity to devaluation of rewards from an action, which naturally maps to economic notions of inelastic demand. We can show that price elasticities will be zero for small price changes within a "zone of price indifference," cross-price elasticities can have different signs, and own-price elasticity can be positive (violating the Law of Demand).

Furthermore, we predict that people exert less mental effort, misremember their purchase history, and may have an impaired ability to consciously explain their choices during habitual choice. Although these variables could in principle be included in a richer model, they are not explicitly represented in the model presented in this chapter – these are simply empirical conjectures that await empirical testing.

References

Abel, A. B. (1990). Asset prices under habit formation and catching up with the Joneses. American Economic Review, Vol. 80, No. 2, pp. 38–42, May.

Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, *33*(2b), 109–121.

Adams, C., & Dickinson, A. (1981). Actions and habits: Variations in associative representations during instrumental learning. *Information Processing in Animals: Memory Mechanisms*, 143–165.

Ahrens, S., Pirschel, I., & Snower, D. J. (2017). A theory of price adjustment under loss aversion. *Journal of Economic Behavior & Organization*, *134*, 78–95.

Andrews, S., Ellis, D. A., Shaw, H., & Piwek, L. (2015). Beyond self-report: Tools to compare estimated and real-world smartphone use. *PloS One*, *10*(10), e0139004.

Atkin, D. (2013). Trade, tastes, and nutrition in India. *American Economic Review*, *103*(5), 1629–1663.

Beauchamp, J. P., Cesarini, D., Johannesson, M., van der Loos, M. J., Koellinger, P. D., Groenen, P. J., … Christakis, N. A. (2011). Molecular genetics and economics. *Journal of Economic Perspectives*, *25*(4), 57–82.

Becker, G. S. (1996). *Accounting for tastes*. Harvard University Press.

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, *96*(4), 675–700.

Benhabib, J., & Bisin, A. (2005). Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption–saving decisions. *Games and Economic Behavior*, *52*(2), 460–492.

Benjamin, D. J., Heffetz, O., Kimball, M. S., & Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose? *American Economic Review*, *102*(5), 2083–2110.

Biggart, N. W., & Beamish, T. D. (2003). The economic sociology of conventions: Habit, custom, practice, and routine in market order. *Annual Review of Sociology*, *29*(1), 443–464.

Block, H. D., & Marschak, J. (1959). *Random orderings and stochastic theories of response*. Cowles Foundation for Research in Economics, Yale University.

Boldrin, M., Christiano, L. J., & Fisher, J. D. (2001). Habit persistence, asset returns, and the business cycle. *American Economic Review*, *91*(1), 149–166.

Bonfrer, A., Berndt, E. R., & Silk, A. (2006). *Anomalies in estimates of cross-price elasticities for marketing mix models: theory and empirical test*. National Bureau of Economic Research.

Bornstein, R. F. (1989). Exposure and affect: overview and meta-analysis of research, *Psychological Bulletin* 106 (2), 265. FINISH CITE

Bronnenberg, B. J., Dubé, J.-P. H., & Gentzkow, M. (2012). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, *102*(6), 2472–2508.

Bucher, T., Collins, C., Rollo, M. E., McCaffrey, T. A., De Vlieger, N., Van der Bend, D., … Perez-Cueto, F. J. (2016). Nudging consumers towards healthier choices: a systematic review of positional influences on food choice. *British Journal of Nutrition*, *115*(12), 2252–2263.

Bucklin, R. E., & Gupta, S. (1992). Brand choice, purchase incidence, and segmentation: An integrated modeling approach. *JMR, Journal of Marketing Research*, *29*(2), 201.

Burga, A., Wang, W., Ben-David, E., Wolf, P. C., Ramey, A. M., Verdugo, C., … Kruglyak, L. (2017). A genetic signature of the evolution of loss of flight in the Galapagos cormorant. *Science*, *356*(6341), eaal3345.

Caballero, R. J., & Engel, E. M. (1999). Explaining investment dynamics in US manufacturing: a generalized (S, s) approach. *Econometrica*, *67*(4), 783–826.

Campbell, J. Y., & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, *107*(2), 205–251.

Caplin, A., Dean, M., & Martin, D. (2011). Search and satisficing. *American Economic Review*, *101*(7), 2899–2922.

L Carden, W Wood (2018) . Habit formation and change, Current Opinion in Behavioral Sciences 20, 117-122

Cecchetti, S. G. (1986). The frequency of price adjustment: a study of the newsstand prices of magazines. *Journal of Econometrics*, *31*(3), 255–274.

Charlton, S. G., & Starkey, N. J. (2011). Driving without awareness: The effects of practice and automaticity on attention and driving. *Transportation Research Part F: Traffic Psychology and Behaviour*, *14*(6), 456–471.

Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy*, *98*(3), 519–543.

Cramer, Jan Salomon. "Interaction of income and price in consumer demand." *International Economic Review* (1973): 351-363.

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

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704.

Daw, N. D., Gershman, S., Seymour, B., Dayan, P. & Dolan, R. (2011). Model-based influences on humans' choices and striatal prediction errors. Neuron, 69(6).

Daw, N. D., & O'Doherty, J. (2014). "Mulitple Systems for Value Learning" in Neuroeconomics (edited by Glimcher, P. & Fehr, E.). Academic Press.

De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, *102*(6), 2955–2980.

Willis, D., Reyes, C. (2015, November 10). The Dog Ate My Vote: How Congress Explains Its Absences [text/html]. Retrieved August 7, 2018, from https://www.propublica.org/article/the-dog-ate-my-vote-how-congress-explains-its-absences

Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences and reinforcement learning. *European Journal of Neuroscience*, *35*(7), 1036–1051.

Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Phil. Trans. R. Soc. Lond. B*, *308*(1135), 67–78.

Dickinson, A., & Balleine, B. (1993). Actions and responses: The dual psychology of behaviour.

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

Duhigg, C. (2012). *The power of habit: Why we do what we do in life and business* (Vol. 34). Random House.

Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of Industrial Organization*, *3*, 1967–2072.

Fehr, E., & Rangel, A. (2011). Neuroeconomic Foundations of Economic Choice--Recent Advances. *Journal of Economic Perspectives*, *25*(4), 3–30.

Frydman, C. & Jin, L. (2021). Efficient Coding and Risky Choice. *Quarterly Journal of Economics.*

Glimcher, P. W. (2011). *Foundations of neuroeconomic analysis*. OUP USA.

Goettler, R. L., & Clay, K. (2011). Tariff choice with consumer learning and switching costs. *Journal of Marketing Research*, *48*(4), 633–652.

Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiology of Learning and Memory*, *70*(1–2), 119–136.

Greenwood, R., & Shleifer, A. (2014). Expectations of Returns and Expected Returns. *The Review of Financial Studies*, *27*(3), 714–746. https://doi.org/10.1093/rfs/hht082

Gul, F., & Pesendorfer, W. (2008). The case for mindless economics. *The Foundations of Positive and Normative Economics: A Handbook*, *1*, 3–42.

Han, S., Gupta, S., & Lehmann, D. R. (2001). Consumer price sensitivity and price thresholds14. *Journal of Retailing*, *77*(4), 435–456.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, *103*(7), 2643–2682.

Hardie, B. G., Johnson, E. J., & Fader, P. S. (1993). Modeling loss aversion and reference dependence effects on brand choice. *Marketing Science*, *12*(4), 378–394.

Hebert, B. & Woodford, M. (2021). Rational Inattention When Decisions Take Time. Working Paper.

Hodgson, G. M., & Knudsen, T. (2004). The firm as an interactor: firms as vehicles for habits and routines. *Journal of Evolutionary Economics*, *14*(3), 281–307.

Hong, H., & Shum, M. (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics*, *37*(2), 257–275.

Hortaçsu, A., Madanizadeh, S. A., & Puller, S. L. (2015). *Power to choose? An analysis of consumer inertia in the residential electricity market*. National Bureau of Economic Research.

Kahneman, D., & Krueger, A. B. (2006). Developments in the measurement of subjective well-being. *Journal of Economic Perspectives*, *20*(1), 3–24.

Kalyanaram, G., & Winer, R. S. (1995). Empirical generalizations from reference price research. *Marketing Science*, *14*(3_supplement), G161–G169.

Keane, Michael P. "Modeling heterogeneity and state dependence in consumer choice behavior." Journal of Business & Economic Statistics 15.3 (1997): 310-327.

Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology*, *7*(5), e1002055.

Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: correcting a reductionist bias. *Neuron*, *93*(3), 480–490.

Laibson, David. "A Cue-Theory of Consumption*." *The Quarterly Journal of Economics* 116, no. 1 (February 1, 2001): 81–119.

Lambert-Pandraud, R., & Laurent, G. (2010). Why do older consumers buy older brands? The role of attachment and declining innovativeness. *Journal of Marketing*, *74*(5), 104–121.

Landry, Peter. "Bad habits and the endogenous timing of urges." The Review of Economic Studies 86.2 (2019): 785-806.

Landry, P. & Webb, R. (2021). "Pairwise Normalization: A Neuroeconomic Theory of Multi-Attribute Choice. Journal of Economic Theory, 193.

Larcom, S., Rauch, F., & 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., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, *81*(3), 687–699.

Lorenz, K. (1935). Der kumpan in der umwelt des vogels. *Journal Für Ornithologie*, *83*(3), 289–413.

Luco, Fernando. "Switching Costs and Competition in Retirement Investment." *American Economic Journal: Microeconomics* 11, no. 2 (May 2019): 26–54.

Luttmer, E. F., & Singhal, M. (2011). Culture, context, and the taste for redistribution. *American Economic Journal: Economic Policy*, *3*(1), 157–179.
Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. MIT Press. *Cambridge, Massachusetts*.
Mayhew, G. E., & Winer, R. S. (1992). An empirical analysis of internal and external reference prices using scanner data. *Journal of Consumer Research*, *19*(1), 62–70.
McFadden, D. (2001). Economic choices. *American Economic Review*, *91*(3), 351–378.
Montgomery, Alan L., and Peter E. Rossi. "Estimating price elasticities with theory-based priors." *Journal of Marketing Research* 36.4 (1999): 413-423.
Moraga‐González, J. L., Sándor, Z., & Wildenbeest, M. R. (2013). Semi‐nonparametric estimation of consumer search costs. *Journal of Applied Econometrics*, *28*(7), 1205–1223.
Nelson, R. R., & Winter, S. (1982). An Evolutionary Theory of Economic Change (Cambridge, Massachusetts and London, Belknap Press of Harvard University Press).
Newman, J. W., & Staelin, R. (1972). Prepurchase information seeking for new cars and major household appliances. *Journal of Marketing Research*, 249–257.
Niv, Y., & Langdon, A. (2016). Reinforcement learning with Marr. *Current Opinion in Behavioral Sciences*, *11*, 67–73.
North, D. C. (1991). Institutions. *Journal of Economic Perspectives*, *5*(1), 97–112.
Nosal, K. (2012). Estimating switching costs for medicare advantage plans. *Unpublished Manuscript, University of Mannheim*.
Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy*, *78*(4, Part 1), 745–763.
Pool, E., Gera, R., Fransen, A., Perez, O., Cremer, A., Aleksic, M., Tanwisuth, S., Quail, S., Ceceli, A., Manfredi, D., Nave, G., Tricomi, E., Balleine, B., Schonberg, T., Schwabe, L., and O'Doherty, J. (2021). Determining the effects of training duration on the behavioral expression of habitual control in humans: a multi-laboratory investigation. PsyArXiv: https://psyarxiv.com/z756h.
Putler, D. S. (1992). Incorporating reference price effects into a theory of consumer choice. *Marketing Science*, *11*(3), 287–309.
Rayo, L., & Becker, G. S. (2007). Habits, peers, and happiness: an evolutionary perspective. *American Economic Review*, *97*(2), 487–491.

Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics*, *53*(8), 1761–1800.

Richards, T. J., Gómez, M. I., & Printezis, I. (2015). Hysteresis, Price Acceptance, and Reference Prices. *American Journal of Agricultural Economics*, *98*(3), 679–706.

Robson, A. J. (2001). Why would nature give individuals utility functions? *Journal of Political Economy*, *109*(4), 900–914.

Robson, A. J. (2002). Evolution and human nature. *Journal of Economic Perspectives*, *16*(2), 89–106.

Ryder, H. E., & Heal, G. M. (1973). Optimal growth with intertemporally dependent preferences. *The Review of Economic Studies*, *40*(1), 1–31.

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

Seger, C. A., & Spiering, B. J. (2011). A critical review of habit learning and the basal ganglia. *Frontiers in Systems Neuroscience*, *5*, 66.

Seiler, S., & Pinna, F. (2017). Estimating search benefits from path-tracking data: measurement and determinants. *Marketing Science*, *36*(4), 565–589.

Sethuraman, R., Srinivasan, V., & Kim, D. (1999). Asymmetric and neighborhood cross-price effects: Some empirical generalizations. *Marketing Science*, *18*(1), 23–41.

Shiller, R. J. (2007). *Low interest rates and high asset prices: An interpretation in terms of changing popular economic models*. National Bureau of Economic Research.

Siegel, S., R. E. Hinson, M. D. Krank, and J. McCully. “Heroin ‘Overdose’ Death: Contribution of Drug-Associated Environmental Cues.” *Science (New York, N.Y.)* 216, no. 4544 (April 23, 1982): 436–37.

Slade, M. E. (1999). Sticky prices in a dynamic oligopoly: An investigation of (s, S) thresholds. *International Journal of Industrial Organization*, *17*(4), 477–511.

Smith, Kyle S, & Graybiel, A. M. (2016). Habit formation. *Dialogues in Clinical Neuroscience*, *18*(1), 33.

Smith, Kyle Stephen, & Graybiel, A. M. (2014). Investigating habits: strategies, technologies and models. *Frontiers in Behavioral Neuroscience*, *8*, 39.

Song, Inseong, and Pradeep K. Chintagunta. "Measuring cross-category price effects with aggregate store data." *Management Science* 52.10 (2006): 1594-1609.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT press.

Tinbergen, N. (1951). The study of instinct.

Webb, R. (2019). The (Neural) Dynamics of Stochastic Choice. Management Science, 65(1).

Wellsjo, Alexandra Steiny. Simple Actions, Complex Habits: Lessons from Hospital Hand Hygiene. UC Berkeley working paper, 2021.

Wesley, M. J., Lohrenz, T., Koffarnus, M. N., McClure, S. M., De La Garza, R., Salas, R., … Montague, P. R. (2014). Choosing money over drugs: the neural underpinnings of difficult choice in chronic cocaine users. *Journal of Addiction*.

Wood, W., & Rünger, D. (2016). Psychology of habit. *Annual Review of Psychology*, *67*, 289–314.

Michael Woodford (2020). Modeling Imprecision in Perception, Valuation, and Choice Annual Review of Economics 12:1, 579-601

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*(6), 464.

Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience*, *19*(1), 181–189.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. Journal of personality and social psychology 9.