How much to purchase? - A cognitive adaptive decision making account

Percy K. Mistry (pkmistry@uci.edu)
Department of Cognitive Sciences, University of California Irvine

Abstract

Repeated purchase decisions often violate assumptions of standard economic or rational choice models, such as demonstrating asymmetric or unstable responses to changes in underlying policy, price, or tax variables. I propose a novel framework for how such decisions can be interpreted through the lens of a cognitive process model. This provides psychologically interpretable characterizations of individuals or population groups. It incorporates mental accounting, hedonic adaptation, confirmation bias, and the influence of perceived trust and fairness. It shows how sequential experiences and contextual aspects such as political affiliation, are mediated by this cognitive process to produce evolving consumption patterns. This novel approach can account for empirically observed violations of conventional choice models. The model is quantitatively fit to experimental data for individual purchase decisions and demonstrates improved descriptive, predictive, and inference capabilities. A proof-of-concept analysis using this model to account for real world consumption trends is also demonstrated.

Introduction

Decisions on what quantity (Q) of a particular item to purchase over time depend on individual preferences, expected benefits from purchasing the item, and associated costs (prices, taxes, etc.). Elasticity (ε) is a canonical economic concept that defines the influence of a unit change of an underlying independent variable (e.g. prices or taxes) on the purchase quantity. Typical choice models assume that elasticities are stable at a population level over relatively long periods of time (controlling for income effects, i.e. the influence of higher levels of income on purchasing power), and that elasticities are symmetric (i.e. respond equally to increases versus decreases). They thus assume that observed changes in purchase quantities (∂Q) in response to changes in underlying variables such as prices and taxes (∂p) can be used to estimate empirical elasticities which in turn can be used to accurately forecast future changes.

Standard choice model: Standard economic models of choice assume that decision makers select the optimal quantity (Q) to purchase by maximizing the net benefit stemming from the utility (U_Q) of owning Q units of an item, less the costs of purchase (pQ), where p is the unit price (equation 1).

$$Q = arg\ max_x \{U_x - px\} \tag{1}$$

Without loss of generalizability, the utility function in equation 2 is assumed for the rest of the paper. This is one of a standard set of utility functions used in behavioral and econometric literature (e.g. Chetty, Looney, and Kroft (2009)). Solving for Q using equations 1 and 2, then taking logs, defining $\varepsilon_p = 1/b$, ($\varepsilon_p > 0$) and $A = -\varepsilon_p\ log(a)$, we obtain equation 3. This is in the form of a log-linear model with a corresponding difference equation 4, with log price elasticity ε_p (a standard economic representation). As per this model $log(Q)$ decreases at a rate of ε_p as $log(p)$ increases and vice versa. Once ε_p is empirically estimated, equation 4 can be used to make forecasts.

$$U_x = \left(\frac{ax^{1-b}}{1-b}\right) \tag{2}$$

$$log(Q) = A - \varepsilon_p\ log(p) \tag{3}$$

$$\partial\ log(Q) = -\ \varepsilon_p\ \partial\ log(p) \tag{4}$$

Evidence against conventional assumptions: There is strong evidence however that elasticities (including, but not limited to, price elasticities such as ε_p described above) may not be stable even in the short run (Hughes, Knittel, & Sperling, 2006; Goodwin, 1992), may not be symmetric (Villas-Boas, Berck, Stevens, & Moe-Lange, 2016; Gately, 1992), may show significant heterogeneity, even directionally, (Chetty, Friedman, Olsen, & Pistaferri, 2009; Ayyagari, Deb, Fletcher, Gallo, & Sindelar, 2009; Fletcher, Frisvold, & Tefft, 2015), and may be easily manipulated by extraneous factors. Whilst these violations are acknowledged, no theory provides a robust and quantitative account of how elasticities evolve over time.

Psychological characterization of dynamic elasticities: In this paper I propose that dynamic characteristics of elasticities can be explained by examining purchase decisions through the lens of a sequential cognitive process. Let $\bar{p}$ define a sequential history of the underlying variable such as prices or taxes, Ψ represent stable cognitive characteristics of an individual or a population (these are elaborated on in subsequent sections), and Δ represent contextual factors (e.g. the measure of political climate or affiliations). Then equation 4 can be replaced with equation 5. Here, $\bar{p}$ and Δ are observable, and cognitive characterization Ψ can be empirically estimated (similar to ε).

$$\partial\ log(Q) = f(\Psi, \bar{p}, \Delta, \partial\ log(p)) \tag{5}$$

Cognitive framework

In this section, I develop a novel cognitive process model to structurally define $f()$ in equation 5, for a sequence of repeat purchase decisions over time. This model is parameterized by psychologically interpretable characteristics Ψ, which interact with sequential history $\bar{p}$ and environmental context Δ to shape continuously evolving patterns of purchases and elasticities. This model is based on bringing together and quantitatively specifying some novel and some previously explored psychological conceptualizations, as elaborated below:

Mental Accounting: Transaction Utility

Thaler (1999, 2008) proposed that consumption quantity (Q) choices were driven by a process of mental accounting that considered a combination of acquisition utility (similar to that defined in equation 2) and transaction utility. Transaction utility reflects the "value" of the deal, typically evaluated against some expectation or reference point, and adds or detracts from the acquisition utility. For this paper, transaction utility is defined by equation 6[1], reflecting the difference between the price p and the expectation or reference price θ.

$$T(Q, p, \theta) = (\theta - p)Q \tag{6}$$

The transaction utility can be positive or negative depending on whether expectations were exceeded. Thus the optimal purchase quantity can now be given by equation 7. This replaces equation 1 of the standard model. The mental accounting theory proposes that the acquisition and transaction utilities are separately evaluated, and may be accorded different weights. Here, δ is a salience weight that emphasizes or reduces the effect of the transaction utility component.

$$Q = arg\ max_x \ \{[U_x - px] + \delta\ [T(x, p, \theta)]\} \tag{7}$$

Salience Weight: Rational, Hedonic, or Altruistic?

The salience weight δ characterizes the nature of decision making. A rational choice would imply $\delta = 0$, since the transaction utility is driven purely by whether or not internal expectations are exceeded, and should not play any role in objective decision making. Applications of the mental accounting framework typically assume $\delta > 0$. This implies that individuals act hedonically in self-interest to maximize the utility they derive from exceeding their internal expectations. In that sense, $\delta > 0$ implies a reference-point that reflects 'the maximum they should be charging me'. Any price lower than this is treated as a positive utility and vice versa. However, some consumption decisions may involve conflicting considerations, such as those of fairness (Xia & Monroe, 2010). For instance, the decision to purchase goods that damage the environment, the decision to evade taxes, or the decision to purchase mandatory health insurance, may result in a conflict between hedonic utility on one hand, and a moral obligation on the other. Such moral obligations can give rise to utilitarian or altruistic concerns (Greene, 2007, 2009) that are concerned with the fairness of policies and redistribution goals. An altruistic reference point may thus reflect 'the bare minimum I should be paying'. For choices involving such moral obligations, if people do indeed demonstrate altruistic concerns, the salience weight may be $\delta < 0$, for at least a non-trivial subset of the population. A price lower than the reference point would reduce transaction utility and vice versa. While this may seem counterintuitive, an example that makes this comprehensible is the case of purchasing goods that are not environmentally friendly. Paying a price higher than the expected reference point may act as a moral justification, and in fact increase the transaction utility and resulting demand by reducing the associated guilt.

Hedonic reference point adaptation

Transaction utilities may be evaluated positive or negatively against a reference point. However this reference point is not typically constant. I propose that the reference point evolves over time (n), motivated by principles of hedonic adaptation (Frederick & Loewenstein, 1999). At time point n the reference point (θ_n) moves closer to the recently experienced values under consideration (e.g. price p_{n-1}), modulated by a hedonic adaptation rate L_h, as shown in equation 8.

$$\theta_n = \theta_{n-1} + L_h(p_{n-1} - \theta_{n-1}) \tag{8}$$

Hedonic adaptation implies that this mechanism serves to increase satisfaction and reduce dissonance created by any large difference between actual prices and expected reference points. This serves to condition people towards recent levels of p. The hedonic adaptation rate L_h may vary by individual or population - higher values of L_h close to 1 imply smaller transaction utility and rational consumption behavior.

Confirmation Bias: Asymmetric adaptation

Confirmation bias, where people place asymmetrical weights on information that confirm rather than contradict their beliefs and actions has been shown to be pervasive over many cognitive processes (Nickerson, 1998; Jones & Sugden, 2001; Palminteri, Lefebvre, Kilford, & Blakemore, 2016). The rate of hedonic adaptation L_h is proposed to be asymmetric, and depends on whether the prospective movement of the reference point supports or inhibits current behavior. Let m reflect a bias that reduces the rate of adaptation ($0 \leq m \leq 1$) when adaptation would serve to inhibit current behavior. This bias is introduced in equation 9. Here, I is an indicator function, with $I = 1$ if ($p < \theta$ under hedonic salience weight $\delta > 0$), or if ($p > \theta$ under altruistic salience weight $\delta < 0$), and 0 otherwise. These situations reflect a inhibition of current behavior based on prospective adaptation, and hence manifest as a lower adaptation rate mL_h. Confirmation bias will thus manifest as a consumption bias, slowing down adaptation that inhibits consumption.

$$\theta_n = \theta_{n-1} + L_h(p_{n-1} - \theta_{n-1})\ (mI + (1 - I)) \tag{9}$$

[1]Note that this assumes a comparison of expectation and realization of the price, however a transaction utility can similarly be expressed based on expectations versus realization for the utility U_x. The framework and model in this paper can be applied without any loss of generalizability to such utility based reference points as well.

Trust based adaptation

Additionally, the reference points are proposed to increase when there is a perception of fairness or trust in the counter-party, and drop otherwise. For example, when considering tax changes and related reference points for tax rates, the government is the counter-party. A reference point for taxes may increase when the government is trusted (e.g. its wealth redistribution goals are considered fair, when political affiliations are in power, and hence higher taxes are more acceptable) than when it is not. Similar considerations may be at work when it comes to prices for goods, and whether people trust a certain brand, or when a brand signals quality, etc. This perception of trust is coded as $\pi = 1$ (trust), $\pi = -1$ (distrust)., or $\pi = 0$ (agnostic). Rate of adaptation in response to these perceptions is governed by L_π, and captured in equation 10. Updating is assumed to occur at every time point n when there is a consumption decision or a change in underlying policy.

$$\theta_n = \theta_{n-1}\,(1+\pi_n L_\pi) + L_h(p_{n-1} - \theta_{n-1})\,(mI + (1-I)) \quad (10)$$

Combined Cognitive-Econometric Model

Here, we replace the standard model in equation 3 by using equations 6, 7, and 10.

Case 1: Purchases Quantity and Price Changes

Equation 7 can now be re-written under the cognitive framework as equation 11, and solved. The log linear demand equation 3 then changes to equation 12. Note that these equations contain the term θ_n given by equation 10[2]. The fully expanded version of equation 12 would thus include the parameters π, L_π, L_h, and m.

$$Q_n = arg\ max_x \left\{ \frac{ax^{1-b}}{1-b} - p_n x + \delta(\theta_n - p_n)x \right\} \quad (11)$$

$$log(Q_n) = A - \varepsilon_p\ log(p_n - \delta(\theta_n - p_n)) \quad (12)$$

Case 2: Purchase Quantity and Tax Changes Next, consider the case where the key variable of interest is how purchase demand may change in response to changes in tax rates t, with the reference point θ being a reference point for what is considered a fair tax rate. There is a lot of evidence to show that there are considerable differences between price and tax elasticity, even when they would have objectively identical impact on consumers (Chetty, Looney, & Kroft, 2009; Chetty, 2015). Following the logic in the previous section, but adding terms for an excise tax rate t that is applied as a percentage on the cost price, so that effective cost would be increased by a value pxt, we obtain equation 13. This considers a situation with constant price p and only changes in the tax rate t and hence a reference point for tax rates only.

$$log(Q_n) = A - \varepsilon_p\ log(p) - \varepsilon_p\ log\big(1 - \delta(\theta_n - t_n)\big) \quad (13)$$

[2]Mathematicaly, extremely high values of the reference point would result in infinite utility, inducing people to spend all resources and maximize the units of consumption. Such reference levels are however *psychologically implausible*, and a mathematical bound for psychological plausibility of θ can be derived in terms of p and δ, such that the $log()$ term in equation 12 never turns negative, implying utility never increases to ∞.

Model Simulation Results

Figure 1 illustrates how different assumptions about the confirmation bias (low or high, governed by m) and mode of processing (hedonic versus altruistic, governed by δ) give rise to systematic deviations from the standard choice model, under different price trend situations (see figure caption for more details). When the predictions from the cognitive model shown in figure 1 are used to infer back what the interpretation of such data would have been under the standard choice model, the resulting inferences reflect highly unstable and variable shifts in conventionally measured elasticity from sub-period to sub-period, as well as asymmetry between elasticity during increasing and decreasing price trends, just as has been reported in literature discussed in the introduction. Such apparent instability and asymmetry is readily explained and generated by stable cognitive characteristics.

Application to Experimental Data (price)

Data: I consider a published experimental dataset from the work reported in Sitzia and Zizzo (2012, 2015). 384 participants made a series of 20 sequential decisions on how many units of a particular lottery to buy. Participants were provided experimental units of currency, and could spend as much of it as they wanted on the lotteries. At the end, the unspent currency, as well as any winnings based on the lotteries were added and converted to real monetary payouts. The lottery remained fixed across all trials, but the purchase price per lottery was varied sequentially. Participants were split into 5 conditions as shown in the left panel in figure 2, where the stimulus (price) patterns for the 5 conditions are represented in different colors. Participants in each of the 5 conditions start with either extremely high (EH), high (H), moderate (M), low (L), or extremely low (EL) levels of prices for the first 10 trials which constitute the "shape" block. All the participants observe the same moderate price levels in the last 10 trials, the "compare" block. The right panel in figure 2 shows the average response (average units bought) for each condition. Key observations made by Sitzia and Zizzo (2012) were that participants with higher initial price purchase more units in the "compare" block than those that have observed a lower initial price. This, as well as the dynamics of many individual patterns of how participants switch purchasing behavior over trials represents a challenge for the standard economic model.

Modeling Results: A standard choice model, as described in the introduction, as well as the cognitive model based on hedonic and asymmetric adaptation (equations 9 and 12) is quantitatively fit to this data, using a Bayesian MCMC framework (JAGS, Plummer et al. (2003)). Since this is an experimental setup, the concept of trust based adaptation is not included in the model. A measure of descriptive fit is evaluated. Additionally, the models are separately tested by providing the first 15 trials for each individual to the models, and obtaining predictions for the last 5 trials. Model comparison using deviance information criteria (DIC, a combined measure of model fit and complexity) was significantly better

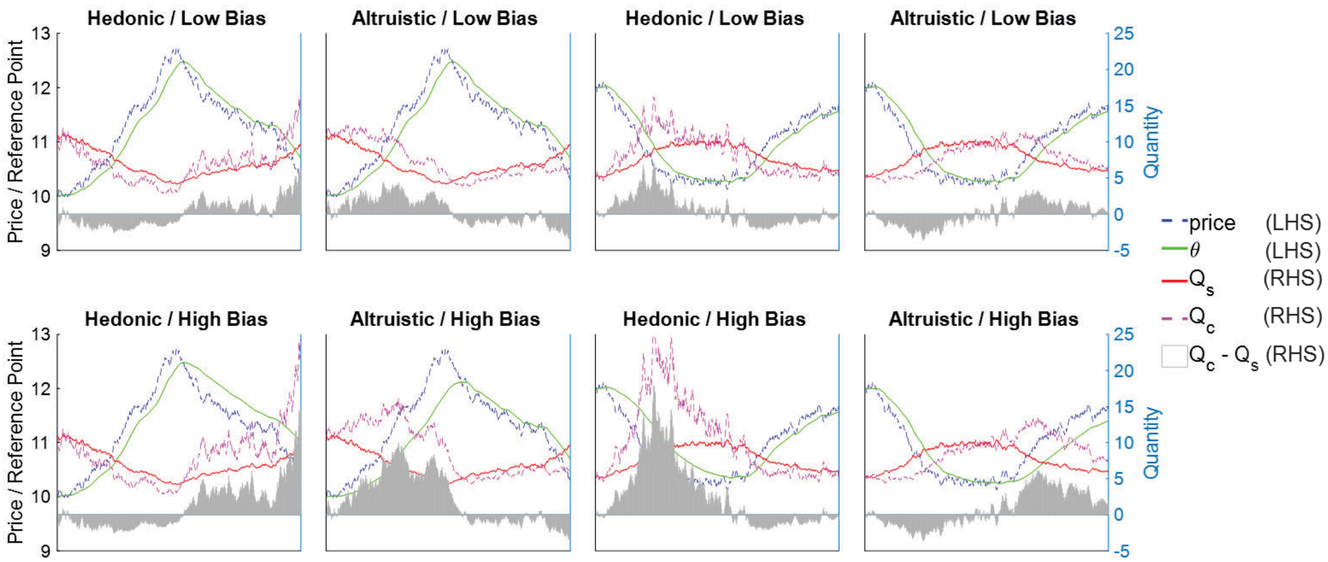

Figure 1: **Model Simulation:** Columns 1 and 2 show price trends that are first increasing and then decreasing. Columns 3 and 4 show price trends that are first decreasing and then increasing. Four parameter combinations reflecting Hedonic ($\delta > 0$) versus Altruistic ($\delta < 0$), and low confirmation bias (high values of *m*) versus low confirmation bias (high values of *m*) are compared under each scenario. The blue lines reflect the price changes. The x-axis reflects time, a hypothetical weekly data spread over a 10 year period. The red line gives the purchase quantity based on the standard choice model and typical assumptions about elasticity. The green line reflects the reference point based on the cognitive model assumptions. The pink line reflects the purchase quantity based on the cognitive model that assumes the same base elasticity as the standard model. The gray bars reflect differences between the cognitive model and the standard model quantities predicted. Price and reference points should be read of the left (LHS) axis and quantities off the right (RHS) axis. High bias parameterizations generally produce higher purchase quantity as expected. More interesting is the asymmetry produced by the cognitive model, which is typically seen in real world scenarios, which can be seen in the relative asymmetry between the gray bars in the first and second half of each simulation - reflecting asymmetries involved in responses to increasing versus decreasing prices.

(lower DIC is better) for the cognitive model (DIC = 21,276) compared to the standard model (DIC = 23,871). Figure 3 compares both the fit and prediction errors (RMSE) between the standard and cognitive models. It shows that for a huge majority of individuals, the cognitive model produces better descriptive fits (better for 86%) and better predictions on unseen data (better for 80%).

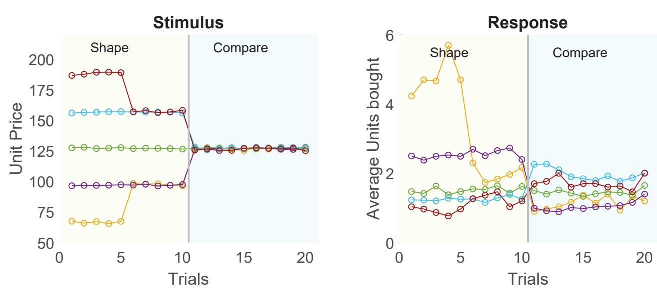

Figure 2: **Stimulus and Responses:** Experimental data from Sitzia and Zizzo (2015)

Illustration of how the model works: Figure 4 shows the latent model inferences about how the reference point evolves, and its relation to consumption patterns for 2 of the participants in the experimental tasks, to illustrate how the model is accounting for behavior. The figure shows the trial by trial stimulus price (red line), the response purchase quantity (black bars) and the latent reference point inferred by the model (green line). For the subject in the left panel, the prices are initially high and then fall. In the second half of the experiment, even though the price stays in the same range, the consumption levels falls as the difference between the latent reference point and the price narrows over time. For the subject in the right panel, the consumption remains almost constant from trials 6-17 even though the price increases after trial 10. This stability when the price is changing significantly is on account of the almost constant difference between the price and the evolving reference point. The standard model finds it difficult to explain these kind of behavioral patterns.

Inferences from the model parameters: Table 1 summarizes the parameter inferences for the cognitive model showing the mean, standard deviation, and the correlation between the parameters and purchase quantities in the "compare" block (trials 11 to 20, where the prices were identical for all 5 conditions). All participants demonstrated con-

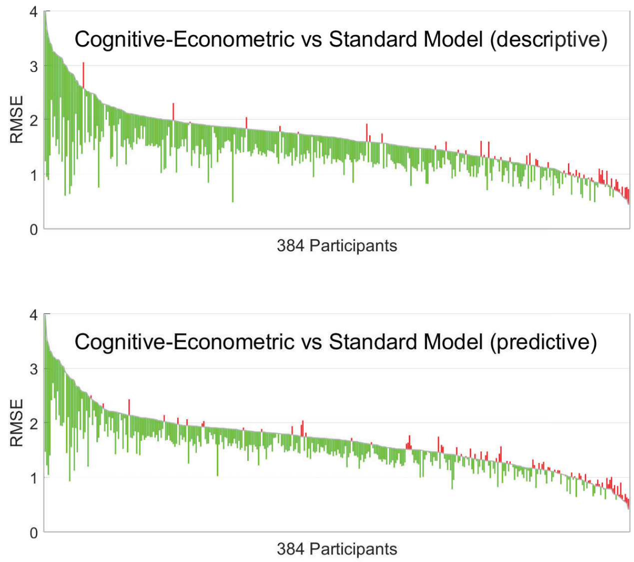

Figure 3: **Comparison of model fit errors:** Significantly better fit (upper panel) and predictions (lower panel) by the cognitive model. Each bar in the figure represents an individual, with the gray line showing the error from the standard model (participants sorted in order of reducing error based on the standard model). The green bars show an improvement (bar going downwards) on account of the cognitive model and red bars show deterioration (bars going upwards).

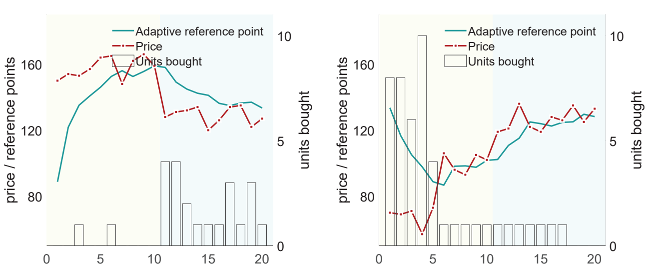

Figure 4: **Illustration of model mechanics:** Examples of the latent model inferences about the reference point trajectory for 2 participants. These behavioral patterns cannot be easily explained by the standard model.

sumption bias ($m < 1$, lower values of m indicate higher consumption bias). The mean value of m of 0.43 indicates that on average, people shifted their reference points twice as much when the price was higher than expectations, than when it was lower than expectations. As intuitively expected, m (lower values = higher consumption bias) is strongly correlated with consumption in the second half of the experiment ($r = -0.75, p < 0.00001$). Most participants show high salience weights δ on the transaction utility, but also individual differences, and higher salience weights are strongly correlated to higher consumption ($r = 0.68, p < 0.00001$). The rate of hedonic adaptation (L_h) did not show strong individual differences, but was consistently less than 1, indicating that adaptation was slower than rationally expected, leading to consistent deviations from the standard model.

Table 1: Key cognitive parameters (Ψ) capturing individual differences in the experimental task.

Characteristics	Mean	std	corr with $Q_{compare}$
m	0.43	0.15	-0.75
δ	1.59	0.71	0.68
L_h	0.48	0.10	0.27

Application to real world data (taxes)

Data: This section provides a brief proof-of-concept for applying this cognitive model to real world population level consumption behavior. Panel data from Chetty, Looney, and Kroft (2009) is used, that includes per capita consumption of beer by state in the US for a period of 34 years, along with the corresponding price and tax changes. As a proof of concept illustration, analysis for 3 states is provided below.

Modeling: A basic standard model[3], and the cognitive model based on equations 10 and 13, that is, including the trust based adaptation, are implemented within a Bayesian inferential framework. The models are fit by providing them with data about consumption changes for 20 years, and then checking model predictions (based on 1000 generated samples for each state for each year) about consumption changes for the last 13 years. The top panels of figure 5 show the changes in tax rates for 3 states over the 34 year period (note the different tax change profiles for the 3 states). The bottom panels show the distribution of prediction errors. The cognitive model produces significantly lower errors ($p < 0.05$ for comparison of error distributions for all 3 states).

Figure 6 shows the influence of the trust based adaptation on reference points (and hence eventually on consumption). The dotted lines show political party regime changes. The three states seem to show graded political affiliations, with the influence of trust switching between high and low (note how the green and blue lines cross over at each regime shift). This is in fact, an inference about the between state differences in trust in existing political regimes, and thus an indicator of state level political affiliation, that was inferred purely from tax rate and beer consumption data. This is an example of the Δ variable suggested in equation 5.

Conclusions

Repeat purchase and consumption decisions are reliant on multiple cognitive processes, and how people respond to

[3]It should be noted that there exist other, more sophisticated and customized econometric models for describing this data. The standard model is used as a baseline comparison to compare the generalizability of standard choice versus cognitive based models.

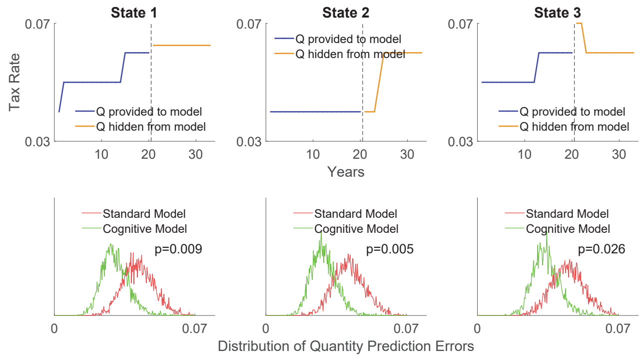

Figure 5: **Application to real world data:** Top panel - Changes in tax rate; Bottom Panel - Prediction error about consumption quantity from standard and cognitive models.

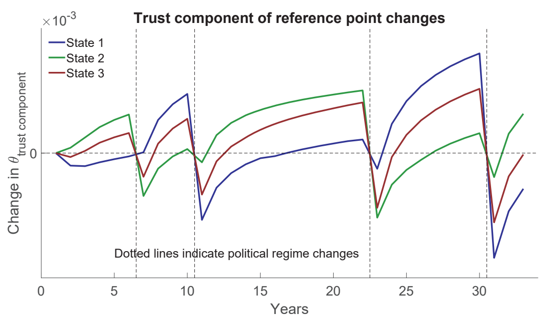

Figure 6: **Effect of Trust based adaptation on reference point:** illustrated are the between state differences based on the shifts in political regime.

changes in prices, taxes, and other policies may deviate significantly from rational models of choice. This paper highlights the importance of a structural model that captures how people's internal expectations may evolve over time, and how capturing this cognitive characterization can help the descriptive and predictive quality of psychological and econometric models. Future work will apply the models to a wider range of experimental and real world data, including identifying heterogeneous sub-population clusters within a larger population (Bell & Lattin, 2000). It will explore the implications for economic predictions, policy implications, and our basic understanding of how adaptive human behavior evolves over time.

References

Ayyagari, P., Deb, P., Fletcher, J., Gallo, W. T., & Sindelar, J. L. (2009). *Sin taxes: do heterogeneous responses undercut their value?* (Tech. Rep.). National Bureau of Economic Research.

Bell, D. R., & Lattin, J. M. (2000). Looking for loss aversion in scanner panel data: The confounding effect of price response heterogeneity. *Marketing Science*, *19*(2), 185–200.

Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *The American Economic Review*, *105*(5), 1–33.

Chetty, R., Friedman, J. N., Olsen, T., & Pistaferri, L. (2009). *Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records* (Tech. Rep.). National Bureau of Economic Research.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *The American economic review*, *99*(4), 1145–1177.

Fletcher, J. M., Frisvold, D. E., & Tefft, N. (2015). Non-linear effects of soda taxes on consumption and weight outcomes. *Health economics*, *24*(5), 566–582.

Frederick, S., & Loewenstein, G. (1999). 16 hedonic adaptation. *Well-being: Foundations of hedonic psychology*, 302.

Gately, D. (1992). Imperfect price-reversibility of us gasoline demand: asymmetric responses to price increases and declines. *The Energy Journal*, 179–207.

Goodwin, P. B. (1992). A review of new demand elasticities with special reference to short and long run effects of price changes. *Journal of transport economics and policy*, 155–169.

Greene, J. D. (2007). Why are vmpfc patients more utilitarian? a dual-process theory of moral judgment explains. *Trends in cognitive sciences*, *11*(8), 322–323.

Greene, J. D. (2009). The cognitive neuroscience of moral judgment. *The cognitive neurosciences*, *4*, 1–48.

Hughes, J. E., Knittel, C. R., & Sperling, D. (2006). *Evidence of a shift in the short-run price elasticity of gasoline demand* (Tech. Rep.). National Bureau of Economic Research.

Jones, M., & Sugden, R. (2001). Positive confirmation bias in the acquisition of information. *Theory and Decision*, *50*(1), 59–99.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, *2*(2), 175.

Palminteri, S., Lefebvre, G., Kilford, E. J., & Blakemore, S.-J. (2016). Confirmation bias in human reinforcement learning: evidence from counterfactual feedback processing. *bioRxiv*, 090654.

Plummer, M., et al. (2003). Jags: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing* (Vol. 124, p. 125).

Sitzia, S., & Zizzo, D. J. (2012). Price lower and then higher or price higher and then lower? *Journal of Economic Psychology*, *33*(6), 1084–1099.

Sitzia, S., & Zizzo, D. J. (2015). Price lower and then higher or price higher and then lower?. [data collection]. uk data service. sn: 851705, http://doi.org/10.5255/ukda-sn-851705.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral decision making*, *12*(3), 183.

Thaler, R. H. (2008). Mental accounting and consumer choice. *Marketing Science*, *27*(1), 15–25.

Villas-Boas, S. B., Berck, P., Stevens, A., & Moe-Lange, J. (2016). Measuring consumer responses to a bottled water tax policy. *American ournal of Ag Economics*.

Xia, L., & Monroe, K. B. (2010). Is a good deal always fair? examining the concepts of transaction value and price fairness. *Journal of Economic Psychology*, *31*(6), 884–894.