Habits as Adaptations: An Experimental Study*

Ludmila Matysková†
University of Bonn

Brian Rogers‡
Washington University in St. Louis

Jakub Steiner§
University of Zurich and CERGE-EI

Keh-Kuan Sun¶
Washington University in St. Louis

September 15, 2018

Abstract

Psychologists emphasize two aspects of habit formation: (i) habits arise when the history of a decision process correlates with optimal continuation actions, and (ii) habits alleviate cognition costs. We ask whether serial correlation of optimal actions alone mechanically induces habit formation, or rather, instead, habits are optimal adaptations. We compare lab treatments that differ in the information that is provided to subjects, holding fixed the serial correlation between optimal actions. We find that past actions affect continuation behavior only in the treatment in which this form of habit is useful. Our result suggests, therefore, that caution is warranted when modeling habits via a fixed utility over action sequences.

Keywords: habit formation, rational inattention.

JEL codes: C91, D8, D9

1 Introduction

Habits play an important role in economic discourse. Economists employ them to explain and predict diverse phenomena ranging from inertia of consumption levels in macroeconomics to brand loyalty in demand estimations. The typical modeling approach represent habits via a fixed time-nonseparable utility function, thus leaving the issues of when and why habits form, and their responses to counterfactual policies, unanswered. Psychologists offer a view on both the purpose and the mechanism of habit formation. They define habits as automated responses triggered

*This paper builds on a dissertation chapter of Ludmila Matyskova entitled "Habit Formation: An Experimental Study". We thank Mark Dean, Emir Kamenica, and Colin Stewart for comments. This work was financially supported by the Czech Science Foundation grant 16-00703S and the ERC grant 770652.

†email: ludmila.matyskova@cerge-ei.cz
‡email: brogers@wustl.edu
§email: jakub.steiner@cerge-ei.cz
¶email: sun.k@wustl.edu

by cues, where cues are elements of the decision history that empirically correlate with optimal continuation choices. In this view the purpose of habits is to alleviate cognition costs.[1]

We ask whether habits originate by mechanically following a cue that empirically correlates with well-performing choices or whether, instead, habits are outcomes of second-best adaptations. Our main experimental result supports the latter hypothesis. We view the finding as good news for the ability to predict habits. Models of habit formation rooted in optimization are able to inform analysts which cues, out of several available cues, a decision maker will adopt. Such models are capable of providing counterfactual predictions of which habits will form under various policies.

To discriminate between the mechanical and optimization origins of habits, we run two treatments of a lab experiment in which subjects face a sequence of decision tasks generated by a fixed stochastic process. The treatments differ only in the information feedback. If habits form mechanically whenever past cues and optimal continuation choices are correlated, then the variation in feedback should not impact habit formation and the selection of cues. Our data, however, show that the subjects form distinct habits across these treatments; moreover, the habits formed can be naturally rationalized as optimal adaptations to the feedback provided.

Our experiment works as follows. The basic task is to recognize a binary state variable presented visually on a computer screen. This state recognition requires moderate cognitive effort. Subjects face two periods of this state-recognition task, across which the state evolves according to a known stochastic process with a positive serial correlation. In the treatment without feedback, we reveal both realized states to the subjects at the end of the second period. In this treatment, we find that subjects form a habit: the first period outcome predicts the second period choice (controlling for the second-period state). The cue that subjects leverage is the first-period *action*; the first-period state is insignificant. In other words, the behavioral pattern exhibits action inertia. The habit alleviates cognitive burden since, due to the serial correlation of the states, the first-period action contains useful information about the second-period optimal choice, and the subjects utilize this information.

In the other treatment, with information feedback, we employ the same state-generating process, but we reveal the first-period state immediately after the first-period action, before the second period begins. Subjects again form a habit in this treatment; payoff-irrelevant elements of the history predict the continuation choice (controlling for the second-period state). However, importantly, the cue changes relative to the previous treatment. The first-period action now becomes insignificant and all of the predictive power is associated with the first-period state, which contains superior information about the optimal continuation action relative to the first-period action. This finding is inconsistent with the view that habits arise as a mechanistic consequence of the serial correlation of optimal actions. Rather, the result suggests that our subjects choose cues optimally, according to their informational content. As a further check, we also ran two additional treatments (with and without the information feedback) in which the states were serially independent. As expected,

[1] See, e.g. Andrews (1903), Lally, Van Jaarsveld, Potts, and Wardle (2010), Wood and Neal (2007) for the psychological studies of habit formation.

subjects do not form habits in those treatments; the second-period choice is independent of all first-period variables.

We supplement the experiment with a model that derives habit formation from primitive assumptions on the information-processing friction. The model is a special case of the dynamic rational-inattention problem from Steiner, Stewart, and Matějka (2017) applied to our experimental setting. In the model, a decision-maker chooses information structures and trades off the precision of her choice against a cost of acquiring information. The model delivers predictions about habit formation consistent with our experimental results. Towards the end of the paper, we discuss the observation that such a model provides counterfactual predictions for habit formation disciplined by optimization arguments in a general setting.

Popular macroeconomic models explain the empirically observed inertia of consumption levels by assuming a time-nonseparable utility function $u\left(c_t - c^{t-1}\right)$, where c^{t-1} is an aggregate of the consumption history, e.g. Pollak (1970) and Abel (1990). When u is concave, a high aggregate past consumption triggers the choice of a high current consumption; i.e., c^{t-1} becomes the cue for a consumption habit. Since the assumed utility representation is exogenous, the modeling choice of c^{t-1} is not obvious and specifications in the literature range from aggregates of past population-wide consumption, individual consumption, or past individual consumption of specific categories of goods; see Schmitt-Grohé and Uribe (2007) for the review of the macroeconomic literature.

Laibson (2001) proposes a model of habit formation rooted in psychology that, like us, focuses on the endogenous selection among several available habit cues, albeit, unlike in our case, the cue selection is not rooted in the optimization of cognition costs. Camerer, Landry, and Webb () study a model of habit formation inspired by neuroeconomics and advocate for the optimization-based origin of habits. Angeletos and Huo (2018) prove observational equivalence between a macroeconomic model featuring higher-order uncertainty, and a model with a representative agent who forms consumption habits.

Our model of habit formation belongs to the rational-inattention literature originating in Sims (2003). As mentioned, it is a special case of the the discrete dynamic rational-inattention model by Steiner, Stewart, and Matějka (2017), which in turn extends a related static model by Matějka and McKay (2015). Rational inattention models have been used to derive inertia of behavior in a macroeconomic context, see e.g. Mackowiak and Wiederholt (2009) for theoretical contribution and Khaw, Stevens, and Woodford (2017) for an experimental exploration.

2 Habits and cues

We study habit formation in the simplest possible setting. A decision-maker (DM) chooses a binary action $a_t \in \{0, 1\}$ in each of two periods to maximize $\sum_{t=1}^{2} u(a_t, \theta_t)$. The binary state $\theta_t \in \{0, 1\}$ evolves according to a stochastic process known to the DM. The DM's task in each period is to match the action to the realized state; $u(a, \theta) = 1$ if $a = \theta$ and zero otherwise.

An analyst collects data on the states and choices in many repetitions of the DM's two-period task. In its idealized form, the analyst observes the joint probability distribution $p(\theta_1, a_1, \theta_2, a_2)$ over the quadruples of the states and the chosen actions. By allowing the analyst to observe the state and action evolution, our data extends the state-dependent stochastic-choice data introduced by Caplin and Dean () in a static setting to the dynamic context considered here.

We say that the DM *forms a habit* if there exists a triple $(\theta_1, \theta_2, a_1) \in \{0, 1\}^3$ such that $p(a_2 \mid \theta_2, \theta_1, a_1) \neq p(a_2 \mid \theta_2)$. Otherwise, if a_2 is independent of (θ_1, a_1) conditional on θ_2, we say that the DM *does not form a habit*. Thus, the DM forms a habit if the history of the process – which is irrelevant for the continuation payoff – predicts continuation behavior. Our definition of habit is behavioral in nature and distinct from its formalization via the non-separable utility approach commonly used in macro-economics. Our analyst knows that the DM's utility is time-separable; she attributes any correlation between the history and continuation behavior to imperfections in the DM's decision process, and refers to the predictive power of the history as to a habit. Implicitly, our definition is related to the concept from Camerer, Landry, and Webb () who define habits as a lack of adaptation to evolving incentives. When, as in our definition, the history of a decision process predicts the continuation behavior controlling for the continuation state, then the DM has not fully adapted to the evolving state.

Additionally, we define cues that trigger the habitual behavior. Is the habitual behavior in the second period, if it arises, triggered by the past state θ_1, or by the past action a_1? Let z be one of the two random variables in the set $\{\theta_1, a_1\}$ and w be the complementary random variable from the same set. We say that *z is the cue for the habit* if (i) $p(a_2 = 1 \mid \theta_2, z = 1) > p(a_2 = 1 \mid \theta_2, z = 0)$, and if (ii) $p(a_2 = 1 \mid \theta_2, z, w) = p(a_2 = 1 \mid \theta_2, z)$. Thus, for instance, the past action a_1 is the cue for the habit if the probability that the DM chooses the high action in the second period increases with a_1 given the second-period state, and the first-period state has no additional predictive power. The latter condition prevents a spurious identification of cues. Since θ_1 and a_1 may be correlated (and indeed they are strongly correlated in our data), it may happen that they both predict choice in the continuation process, but all the predictive information is contained in only one of them.

Our experiments vary the two-period decision problem along two dimensions: (i) we consider independent or positively serially correlated states θ_t, and (ii) we reveal or do not reveal the realization of the first-period state to the DM before the second period begins. Based on the hypothesis that habits are useful adaptations, we predict that the DM does not form a habit when the states are independent (for either information treatment). When states are positively correlated and the first-period state is not revealed then we predict that the DM forms a habit with the cue a_1. In this case, the first-period action contains useful information about the optimal second-period choice, and the DM may reduce her cognitive costs by partially relying on this information. When states are correlated and the first-period state is revealed then we predict that the DM forms a habit with the cue θ_1, since θ_1 contains information about θ_2 superior to information contained in a_1.

In the next section, we confirm the above intuitive predictions in a formal model, after which we proceed with a description of the lab experiment, followed by a summary of our empirical results.

3 Rational-inattention model of habits

In this section, we briefly review a rational-inattention model of habit formation. It deviates from the standard macroeconomic modeling of habits in that it keeps the DM's utility fixed and time-separable; it explains habits as optimal second-best adaptations to an information-processing friction. The model is a special case of Steiner, Stewart, and Matějka (2017).

The DM conducts a costly statistical experiment that produces a signal x_t in each period $t = 1, 2$. Additionally, in between the periods 1 and 2, she receives an exogenous signal y produced by a process beyond the DM's control. In each period t, she chooses an action according to a (pure) action strategy σ_t that maps the observed signals up to period t to a_t; that is, $a_1 = \sigma_1(x_1)$ and $a_2 = \sigma_2(x_1, x_2, y)$. The DM controls the statistical experiments that generate signals x_t and can condition the choice of the employed experiment on all the available information at the given period. Let X, $|X| \geq 2$, be a fixed signal space. The DM chooses any statistical experiment $f_1(x_1 \mid \theta_1)$ and any system of statistical experiments $f_2(x_2 \mid \theta_2, x_1, y)$ that govern the conditional probability distribution of the signals $x_t \in X$ for each combination of the values of the random variables specified in the condition.[2]

We consider two distinct processes that generate the exogenous signal y. In one case, the exogenous signal perfectly reveals the first state; $y = \theta_1$. We say in this case that the *DM receives feedback*. In the other case, $y = y_0$, where y_0 is an arbitrary constant, and we say that the *DM does not receive feedback*.

The DM chooses the statistical experiments and the action strategies to maximize her expected payoff net of the information cost. That is, she solves

$$\max_{f_1, f_2, \sigma_1, \sigma_2} \mathrm{E}\left[u\left(\sigma_1\left(x_1\right), \theta_1\right) + u\left(\sigma_2\left(x_1, x_2, y\right), \theta_2\right) - \lambda\left(I\left(\theta_1; x_1\right) + I\left(\theta_2; x_2 \mid x_1, y\right)\right)\right], \tag{1}$$

where the expectation is over the random variables θ_1, θ_2, x_1, x_2, and y.[3] The DM's cost of the first signal x_1 is proportional to the *mutual information* $I(\theta_1; x_1)$ that measures the informativeness of the signal x_1 about the random variable of the DM's interest, θ_1. The cost of the second signal is proportional to the *conditional mutual information* $I(\theta_2; x_2 \mid x_1, y)$ that measures informativeness of x_2 about the state θ_2 relative to the information contained in x_1 and y. See the Appendix for formal definitions. The parameter $\lambda > 0$ scales the information costs.

Let $p^*(\theta_1, a_1, \theta_2, a_2)$ be the joint distribution of the states and actions generated by the optimal experiments $f_1^*(x_1 \mid \theta_1)$ and $f_2^*(x_2 \mid \theta_2, x_1, y)$ and action strategies $\sigma_1^*(x_1)$, $\sigma_2^*(x_1, x_2, y)$. We impose the regularity conditions that (i) θ_1 and θ_2 attain both values with positive probabilities, (ii) θ_1 and a_1 are correlated, and (iii) a_2 attains both values with positive probabilities.[4]

The next proposition confirms the intuition from the end of the previous section.

[2] The signal x_1 is constrained to be independent from θ_2 conditional on θ_1. Similarly, x_2 is constrained to be independent from θ_1 conditional on (θ_2, x_1, y); the DM learns about θ_t only in period t.

[3] The joint probability distribution of these variables is fully specified by the stochastic process of θ_t, by the chosen statistical experiments f_t and by the process that generates the exogenous signal y.

[4] Parts (ii) and (iii) must hold when λ is sufficiently low, and they hold in our experimental data.

Proposition 1. *Suppose that the evolution of the decision process is described by p^*. Then:*

1. *If θ_1 and θ_2 are independent, then the DM does not form a habit.*
2. *If θ_1 and θ_2 are positively correlated and the DM does not receive feedback, then she forms a habit with the cue a_1.*
3. *If θ_1 and θ_2 are positively correlated and the DM receives feedback, then she forms a habit with the cue θ_1.*

4 Experimental design

Our design follows Caplin and Dean (2015). Subjects are presented with an image of a ten by ten matrix of red and blue balls on a computer screen. There are two possibilities (states): either 51 red and 49 blue or 51 blue and 49 red balls are displayed. The position of the balls is random; see the screenshot in Figure 1 in the Appendix. The subjects are incentivized to determine the majority color. They do not face any explicit cost of gathering information and can perfectly learn the realized state; information costs stem from real cognitive effort. We refer to the above one-period decision problem as the *counting task*. For a technical reason, a 45 second time limit is imposed after which the image of the balls disappears and a decision is forced.[5]

Since we are interested in serial correlations that arise in the absence of real switching costs, we let the position of the buttons (labeled Red and Blue) via which the subjects indicated their choice randomly alternate across periods. Thus, provision of the same answer in consecutive periods was not associated with a mental or physiological advantage.

We recruited 41 subjects from the University of California, Santa Barbara over 2 sessions. In each session, subjects faced 4 treatments. Each treatment consisted of 12 iterations and each iteration consisted of 2 periods of the counting tasks. Thus, each subject faced 96 counting tasks in total. In each iteration, both state realizations were equally likely in the first period. The four treatments per session are defined by the combinations of: (i) the serial correlation between the states, and (ii) informational feedback in between the two periods.

First, in the *independent* treatments (I), the states θ_1 and θ_2 drawn from the set $\{\text{blue, red}\}$ were independent. In *correlated* treatments (C), the states were positively correlated and $\Pr(\theta_1 = \theta_2) = 3/4$. The pairs of states were drawn from a given distribution by a computer. Second, in the *feedback* treatments (F), the subjects were shown the realized state θ_t immediately after choosing a_t. Hence, the subjects perfectly learned the realization of the first state before entering the second period, independently of their learning effort. In the *no feedback* treatments (N), the subjects were shown the realized pair of states only after actions in both periods were taken. The four treatments are captured in Table 1. Prior to starting each treatment, subjects were informed about the state-generating process and the feedback specification. The order of the treatments was: FI, FC, NI,

[5] The time limit was set to ensure a reasonable duration for the experimental sessions, yet such that the subjects do not feel pressured by the time limit.

Table 1: Treatments

	feedback	no feedback
independent states, $\Pr(\theta_1 = \theta_2) = \frac{1}{2}$	FI	NI
correlated states, $\Pr(\theta_1 = \theta_2) = \frac{3}{4}$	FC	NC

NC in Session 1 and NI, NC, FI, FC in Session 2. Within a treatment, each subject faced the same sequence of images.

Subjects were informed that their payoff depends on their choices and chance as follows. At the end of the experiment, for each subject one counting task was selected uniformly at random from all 96 tasks and the payment based on the subject's answer to the selected counting task was added to a \$10 show-up fee. The marginal payment was \$10 for a correct answer and \$0 otherwise. Hence, the expected additional payoff for a correct answer in any given task was $\$\frac{10}{96} \approx \$0.10$.[6] See the Appendix for our experimental instructions.

5 Experimental results

Since in each treatment each of 41 subjects faced 12 repetitions of the two-period problem, we have 492 observations of the quadruple $(\theta_1, a_1, \theta_2, a_2)$ for each treatment. The precision of choice was high and similar across treatments. The proportion of correct choices pooled across subjects in each period of each treatment ranged from 84% to 90%. Subjects exhibited serial correlations of their actions in the correlated treatments. The empirical frequency that $a_1 = a_2$ was 0.78 in both the FC and NC treatments.[7] Subjects displayed heterogeneous abilities; the number of correctly answered tasks per subject varied from 51 to 96, with a mean of 83 and a standard deviation of 13.5.

The empirical frequency that $\theta_2 = a_1$ was 0.77 in the FC treatment, 0.74 in the NC treatment, 0.50 in the NI treatment, and 0.61 in the FI treatment. This value summarizes how predictive a_1 is of θ_2. A t-test of whether the empirical frequency $\Pr_{FC}(\theta_2 = a_1)$ significantly differs from $\Pr_{NC}(\theta_2 = a_1)$ is rejected at a 5% significance level.

Our main interest is to examine how the first-period state θ_1 and action a_1 predict the second-period action a_2 in four logit regressions separately for each treatment:

$$a^n_{2,i} = \begin{cases} 1 & \text{if } \beta_0 + \beta_{a_1} a^n_{1,i} + \beta_{\theta_1}\theta^n_1 + \beta_{\theta_2}\theta^n_2 + \beta_{session_2} session_2 + \beta_{score} score^n_i \theta^n_2 + \varepsilon^n_i > 0, \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

[6] The provided incentives are comparable to those in Caplin and Dean (2013), where the expected payoff per correct answer in each counting task varied from \$0.01 to \$0.15.

[7] The empirical probability that $a_1 = a_2$ was .51 in the NI treatment, and .60 in the FI treatment. In the latter case, the realized frequency of $\theta_1 = \theta_2$ was 0.67, and thus the observed action inertia is due to the subjects' attentiveness to the state realizations.

Table 2: Average marginal effects and standard errors in regressions (2) for each treatment.

	a_1	θ_1	θ_2	$session_2$	$score$
FI	-.021 (.036)	.071 (.043)	**.681***** (**.032**)	-.005 (.027)	**.002*** (**.001**)
NI	.034 (.041)	-.026 (.049)	**.692***** (**.054**)	-.011 (.035)	**.004***** (**.001**)
FC	.017 (.032)	**.258***** (**.058**)	**.611***** (**.046**)	-.038 (.019)	.001 (.001)
NC	**.191***** (**.051**)	.002 (.036)	**.629***** (**.067**)	.026 (.026)	**.004***** (**.001**)

with robust standard errors clustered at the subject level, where $a_{t,i}^n$ is the action taken by subject $i = 1, \ldots, 41$ in iteration $n = 1, \ldots, 12$ and period $t = 1, 2$; θ_t^n is the realized state in iteration n and period t; $session_2$ is a dummy variable indicating session 2. Finally, $score_i^n$ is a subject-specific proxy for counting ability. It is the total number of correct answers by subject i in all treatments excluding answers in iteration n of the considered treatment (to avoid endogeneity). The term $score_i^n \theta_2^n$ in the regression captures the idiosyncratic sensitivity of the subject to the variation in second-period state. Not accounting for subjects' heterogeneity could result in an omitted-variable bias.

Table 2 reports the estimated average marginal effects and standard errors of the explanatory variables; see the Appendix for regression outputs.

Experimental results. 1. *Subjects pay attention to the second-period state:*

In all four treatments, variation in a_2 is predominantly explained by the variation in θ_2.

2. *When the states are independent, subjects do not form habits:*

   In treatments FI and NI, θ_2 is the only statistically significant explanatory variable (apart from the score).

3. *When the states are positively correlated and the subjects do not receive feedback, they form a habit with cue a_1:*

   In treatment NC, θ_2 and a_1 are positively statistically significant and θ_1 is insignificant.

4. *When the states are positively correlated and the subjects receive feedback, they form habit with cue θ_1:*

   In treatment FC, θ_2 and θ_1 are positively statistically significant and a_1 is insignificant.

6 Discussion

Relative to models that assert mechanistic habit formation or that represent habits via fixed time-nonseparable utilities, models that derive habits from optimization regarding information processing have superior predictive power. The optimization-based models of habits may predict how cues are selected, and how the nature and strength of the habits change with the DM's environment. The dynamic rational-inattention model of Steiner, Stewart, and Matějka (2017) provides a particularly simple optimization-based representation of habits. Suppose that the DM maximizes expected time-separable utility $\sum_t^T u_t(a_t, \theta_t)$ less the entropy-based information cost as in Section 3. Let the stochastic choice function $p(a_t \mid \theta_t, h^{t-1})$ specify the probability that the DM chooses action a_t in period t when the current state is θ_t and the history of the decision process is $h^{t-1} = (a_1, y_1 \ldots, a_{t-1}, y_{t-1})$, where $y_{t'}$ is an exogenous signal that the DM has received in period t'. Steiner, Stewart, and Matějka (2017) show that the stochastic choice function of the rationally inattentive DM coincides with the dynamic logit choice function of Rust (1987) with *modified* utility $\tilde{u}_t(a_t, h^{t-1}, \theta_t) = u_t(a_t, \theta_t) + \lambda \log p(a_t \mid h^{t-1})$. The latter term is the (endogenous) conditional probability of the action a_t at the history h^{t-1}. This is a quasi-utility term that represents the habitual behavior in favor of the action a_t at history h^{t-1}. The smaller $p(a_t \mid h^{t-1})$ is, the more surprising the choice of a_t is at action history h^{t-1}, and the larger is the quasi-utility penalty $\lambda \log p(a_t \mid h^{t-1})$ associated with this choice. Thus, for a fixed environment, this particular model represents habits via quasi-(dis)utilities of surprising choices.

The above model predicts counterfactual comparative statics of the quasi-utility terms that represent habits. When a policy affects incentives or the stochastic process that governs the state evolution, the endogenous conditional action probabilities $p(a_t \mid h^{t-1})$ change, and hence the quasi-utility habit-representing terms adjust. If, for instance, a policy subsidizes an action a_t at a history h^{t-1}, then a_t becomes more commonly chosen at h^{t-1} and thus its choice at this history becomes more habitual.

In summary, if habits are, as our experimental findings suggest, an outcome of optimization, then economic models of habits may improve their counterfactual predictive power by deriving habits from optimization in micro-founded models.

References

Abel, A. (1990). Asset prices under habit formation and catching up with the joneses. *American Economic Review 80*(2), 38–42.

Andrews, B. R. (1903). Habit. *The American Journal of Psychology 14*(2), 121–149.

Angeletos, G.-M. and Z. Huo (2018). Collective myopia and habit. In *2018 Meeting Papers*, Number 236. Society for Economic Dynamics.

Camerer, C., P. Landry, and R. Webb. Incorporating habit in to consumer demand theory: A neuroeconomic approach. Technical report.

Caplin, A. and M. Dean. Revealed preference, rational inattention, and costly information acquisition. Technical report, National Bureau of Economic Research Working Paper 19876.

Caplin, A. and M. Dean (2013). Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research Working Paper 19318.

Caplin, A. and M. Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review 105*(7), 2183–2203.

Khaw, M. W., L. Stevens, and M. Woodford (2017). Discrete adjustment to a changing environment: Experimental evidence. *Journal of Monetary Economics 91*, 88–103.

Laibson, D. (2001). A cue-theory of consumption. *The Quarterly Journal of Economics 116*(1), 81–119.

Lally, P., C. H. Van Jaarsveld, H. W. Potts, and J. Wardle (2010). How are habits formed: Modelling habit formation in the real world. *European Journal of Social Psychology 40*(6), 998–1009.

Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review 99*(3), 769–803.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–98.

Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy 78*(4, Part 1), 745–763.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica 55*(5), 999–1033.

Schmitt-Grohé, S. and M. Uribe (2007). Habit persistence. *The New Palgrave Dictionary of Economics, second edition, Palgrave Macmillan, NY, NY*.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Steiner, J., C. Stewart, and F. Matějka (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica 85*(2), 521–553.

Wood, W. and D. T. Neal (2007). A new look at habits and the habit-goal interface. *Psychological Review 114*(4), 843.

A Appendix

Definitions:

For a random variable W with finite support S distributed according to $q \in \Delta(S)$, the *Shannon entropy* of W is defined as

$$H(W) = -\sum_{w \in S} q(w) \log q(w),$$

Figure 1: A representative screenshot of a counting task

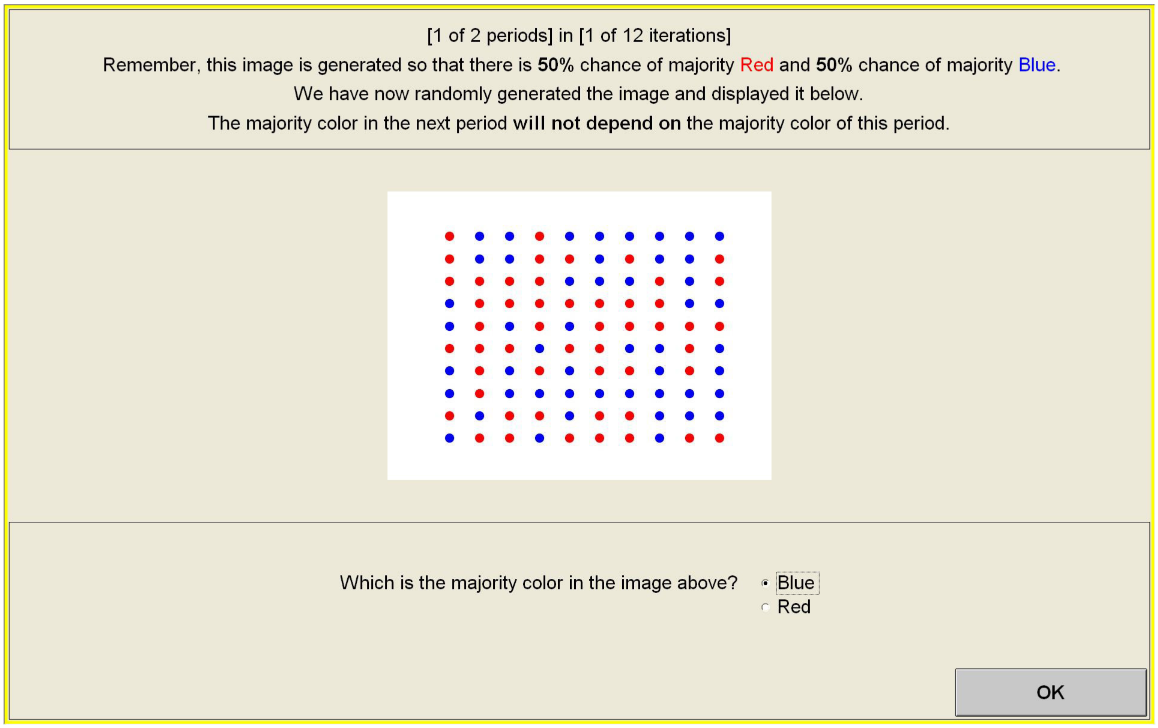

where $0 \log 0 = 0$ by convention. The *mutual information* $I(X;Y)$ of two random variables, X and Y, with finite supports is defined as

$$I(X;Y) = H(X) - \mathrm{E}_Y[H(X \mid Y)].$$

Finally, the *conditional mutual information* $I(X;Y \mid Z)$ is defined as

$$I(X;Y \mid Z) = H(X \mid Z) - \mathrm{E}_Y[H(X \mid Y, Z)].$$

Proof of Proposition 1. statement 1: Signals x_1 and y are independent of θ_2. The second-period continuation problem is the static rational-inattention problem from Matějka and McKay (2015) that is independent of values of θ_1 and x_1. Since this static problem has a unique solution that specifies the stochastic choice rule $p(a_2 \mid \theta_2)$, a_2 must be, conditionally on θ_2, independent of θ_1 and x_1, and hence of $a_1 = \sigma_1(x_1)$ too.

statement 2: Steiner, Stewart, and Matějka (2017) show that the DM chooses experiments f_1 and f_2 that attain only two signals values, the two signal values can be labeled without loss of generality by 0 and 1, and the DM uses action strategies $\sigma_1^*(x_1) = x_1$ and $\sigma_2^*(x_1, x_2) = x_2$. Since the uninformative signal y attains a single value y_0, we omit it from the arguments of σ_2^* and f_2^*. Thus, we have that $p(a_2 \mid \theta_2, \theta_1, a_1) = f_2^*(x_2 = a_2 \mid \theta_2, x_1 = a_1) = p(a_2 \mid \theta_2, a_1)$, as needed.

Therefore it suffices to prove that $p(a_2, \mid \theta_2, a_1)$ increases in a_1.

We first show that $a_1 = x_1$ is positively (rather than negatively) correlated with θ_1. Suppose for contradiction that $a_1 = x_1$ is negatively correlated with θ_1. Then, there exists a payoff improving decision process that uses the same statistical experiments $f_1^*(x_1 \mid \theta_1)$ and $f_2^*(x_2 \mid \theta_2, x_1)$ and action strategies $\tilde{\sigma}_1(x_1) = \sigma_1^*(1-x_1)$, $\tilde{\sigma}_2(x_1, x_2) = \sigma_2^*(x_1, x_2)$. This alternative decision process leads to the same information cost as the original process and to the same expected payoff in period 2, but it is payoff improving in period 1. Hence, the original process could not have been optimal, as needed for contradiction. Second, the continuation problem in period 2 is a static rational-inattention problem of Matějka and McKay (2015) with the prior belief over θ_2 equal to the conditional probability distribution $p(\theta_2 \mid a_1)$. Since θ_1 and θ_2 are positively correlated, and a_1 is positively correlated with θ_1, this prior distribution increases in a_1. Proposition 3 in Matějka and McKay (2015) implies that $p(a_2, \mid \theta_2, a_1)$ increases with a_1, and this monotonicity is strict since a_2 attains both values with positive probabilities.

statement 3: The belief in period 2 about θ_2 is independent of $x_1 = a_1$; that is, $p(\theta_2 \mid x_2, x_1, y) = p(\theta_2 \mid x_2, \theta_1 = y)$. Thus, $\sigma_2^*(x_1, x_2, y)$ is independent of $x_1 = a_1$, and hence $p(a_2 \mid \theta_2, \theta_1, a_1) = p(a_2 \mid \theta_2, \theta_1)$, as needed. The continuation problem in period 2 is the static rational-inattention problem with the prior $p(\theta_2 \mid \theta_1)$, which increases in θ_1, since the two states are positively correlated. Again, by Proposition 3 from Matějka and McKay (2015), the choice rule solving this continuation problem increases in θ_1. $\square$

Instructions

Welcome to the experiment! Please take a record sheet at the front if you don't have one already. Please do not use the computers during the instructions. When it is time to use the computer, please follow the instructions precisely.(Repeat if necessary.)

Please raise your hand if you need a pencil. Please put away and silence all your personal belongings, especially your phone. We need your full attention during the experiment.

Raise your hand at any point if you cannot see or hear well.

The experiment you will be participating in today is an experiment in decision making. At the end of the experiment, you will be paid for your participation in cash. The amount you earn depends on your decisions and on chance. You will be using the computer for the experiment, and all decisions will be made through the computer. DO NOT socialize or talk during the experiment.

All instructions and descriptions that you will be given in the experiment are accurate and true. In accordance with the policy of this lab, at no point will we attempt to deceive you in any way.

If you have any questions, raise your hand and your question will be answered out loud so everyone can hear.

After you have completed all the tasks, please wait while everyone else finishes his or her tasks. Once everyone has completed the experiment, I will ask you to fill in the questionnaire. After the questionnaire you will collect your earnings and leave.

//

I will now describe the main features of the experiment and show you how to use the software. Again, if you have any questions during this period, please raise your hand.

You will be presented with a series of choices to make. There will be four SETS of choices in today's experiment. Each set contains twelve ITERATIONS, and each iteration has two PERIODS. In each period, you will be shown a picture of 100 dots. Each dot will be either RED or BLUE. We have displayed an example of such a screen on your computer monitor. (show an example screen)

This is an example of the screens you will see during the experiment. In every period, the picture will contain either 51 red dots and 49 blue dots, or instead, 51 blue dots and 49 red dots. We will call these two cases MAJORITY RED and MAJORITY BLUE, respectively. In each case, the dots are randomly allocated to the positions in the matrix. In each period the computer will choose randomly between MAJORITY RED and MAJORITY BLUE. You will be told in advance how likely each case is to happen.

In each period, you will be asked to determine if the image is MAJORITY RED or MAJORITY BLUE. While you may take as much time as you need to make your choice, the image will disappear after 45 seconds.

I am now going to describe the details of the experiment.

The experiment is divided into four SETS. In each set, you will be presented with twelve iterations, and each iteration consists of two periods, each with its own image. The rules for the 12 iterations within each set are identical, but the rules are different in different sets.

In PERIOD 1 of each iteration, the image is always generated so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, meaning that there is a 50% chance of MAJORITY RED and a 50% chance of MAJORITY BLUE.

In period 2 of each iteration, the image will be generated in a way that differs across sets. In some sets, the majority color for period 2 is chosen in a way that is completely separate from the period 1 image, and is randomly generates so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, just like the period 1 image. But in other sets, the period 2 image depends on the majority color of the period 1 image. In these sets, the computer generates the period 2 image so that there is a 75% chance that the majority color matches the period 1 majority color, and a 25% chance that the majority color is different from the period 1 majority color.

It is important to remember that while the periods within each iteration may be related to each other, the periods across iterations are never related.

After making your choices, you will always be told what the majority color was, but the timing of this differs from set to set. In some sets, the majority colors will be revealed after every period. In other sets, the majority colors for an iteration will not be revealed until you complete both periods. Before each set, you will be told about the timing of the feedback you will receive.

The amount of money you will receive at the end of the experiment depends on your choices. After we have completed all four sets, you will have made 96 choices (4 sets times 12 iterations times 2 periods). The computer software will randomly select one of these 96 periods. Your payment will be determined by your choice in that single period. If your choices in the randomly chosen period matches the majority color, you will earn an additional $5 dollars on top of the $15 show-up fee. Otherwise, you will receive no additional payment, but you will still receive the show-up fee.

After you complete the last set, please wait until we start the questionnaire part. After you finish the questionnaire, please fill your record sheet on the desk. I will pay one by one to keep everyone's privacy.

To summarize, remember that we have four sets in the experiment today. Each set consists of 12 iterations, and each iteration consists of two periods. The sets will vary in how likely it is that the majority colors are the same for both periods within an iteration, and in the timing that the majority colors are revealed. Please raise your hand if you have any questions.

(1) FI/FC/NI/NC

Feedback/IID:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, and it does not depend on the majority color in the first period.

The majority colors will be revealed after every period, so that you will be told the majority color from period 1 before you see the image for period 2. Please raise your hand if you have any question.

Feedback/Corr.:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is a 75% chance that the majority color matches the majority color from period 1, and a 25% chance that the majority color is different from period 1.

The majority colors will be revealed after every period, so that you will be told the majority color from period 1 before you see the image for period 2. Please raise your hand if you have any question.

No Feedback/IID:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, and it does not depend on the majority color in the first period.

The majority colors for both periods of an iteration will be revealed only at the end of each iteration, so that you will see the period 2 image before being told the majority color from period 1. Please raise your hand if you have any question.

No Feedback/Corr.:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is a 75% chance that the majority color matches the period 1 majority color, and a 25% chance that the majority color is different from the period 1 majority color.

The majority colors for both periods of an iteration will be revealed only at the end of each iteration, so that you will see the period 2 image before being told the majority color from period 1. Please raise your hand if you have any question.

(2) NI/NC/FI/FC

No Feedback/IID:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, and it does not depend on the majority color in the first period.

The majority colors for both periods of an iteration will be revealed only at the end of each iteration, so that you will see the period 2 image before being told the majority color from period 1. Please raise your hand if you have any question.

No Feedback/Corr.:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is a 75% chance that the majority color matches the period 1 majority color, and a 25% chance that the majority color is different from the period 1 majority color.

The majority colors for both periods of an iteration will be revealed only at the end of each iteration, so that you will see the period 2 image before being told the majority color from period 1. Please raise your hand if you have any question.

Feedback/IID:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is an equal chance of MAJORITY RED and MAJORITY BLUE, and it does not depend on the majority color in the first period.

The majority colors will be revealed after every period, so that you will be told the majority color from period 1 before you see the image for period 2. Please raise your hand if you have any question.

Feedback/Corr.:

In the next set of twelve iterations, the majority color for period 2 is randomly generated so that there is a 75% chance that the majority color matches the majority color from period 1, and a 25% chance that the majority color is different from period 1.

The majority colors will be revealed after every period, so that you will be told the majority color from period 1 before you see the image for period 2. Please raise your hand if you have any question.

Regressions outputs

FI TREATMENT

Logistic regression	Number of obs	=	492
	Wald chi2(6)	=	142.93
	Prob > chi2	=	0.0000
Log pseudolikelihood = -163.10254	Pseudo R2	=	0.5150

(Std. Err. adjusted for 41 clusters in subject)

choice2	Odds Ratio	Robust Std. Err.	z	P>|z|	[95% Conf.	Interval]
1.choice1	.8083733	.2882132	-0.60	0.551	.4019102	1.625904
1.state1	1.96372	.7578438	1.75	0.080	.921683	4.183862
1.state2	.0000563	.0000942	-5.85	0.000	2.12e-06	.0014955
1.session2	.9545169	.2568266	-0.17	0.863	.5633204	1.617379
state2#c.score						
0	.9325278	.0140836	-4.63	0.000	.905329	.9605438
1	1.108449	.0179016	6.38	0.000	1.073912	1.144097
_cons	35.59983	41.01105	3.10	0.002	3.722741	340.4341

```
. margins, dydx(*)
```

Average marginal effects	Number of obs	=	492

Model VCE : Robust

Expression : Pr(choice2), predict()
dy/dx w.r.t. : 1.choice1 1.state1 1.state2 1.session2 score

	dy/dx	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
1.choice1	-.0213619	.0355646	-0.60	0.548	-.0910673	.0483434
1.state1	.0717546	.042604	1.68	0.092	-.0117478	.1552569
1.state2	.680717	.0323603	21.04	0.000	.617292	.7441419
1.session2	-.0046997	.0270881	-0.17	0.862	-.0577913	.048392
score	.0021162	.0009961	2.12	0.034	.000164	.0040684

Note: dy/dx for factor levels is the discrete change from the base level.

NI TREATMENT

```
Logistic regression                               Number of obs   =        492
                                                  Wald chi2(6)    =     133.54
                                                  Prob > chi2     =     0.0000
Log pseudolikelihood = -169.51864                 Pseudo R2       =     0.3798

                                 (Std. Err. adjusted for 41 clusters in subject)
--------------------------------------------------------------------------------
               |               Robust
       choice2 | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
---------------+----------------------------------------------------------------
     1.choice1 |   1.389138   .5456704     0.84   0.403     .6432551    2.999907
      1.state1 |   .7762836   .3740146    -0.53   0.599     .3019328    1.995863
      1.state2 |   .0013902   .0030984    -2.95   0.003     .0000176     .109694
    1.session2 |   .8971487   .3004906    -0.32   0.746     .4653324    1.729679
               |
state2#c.score |
            0  |   .9428402   .0218128    -2.54   0.011     .9010427    .9865765
            1  |   1.073454   .0132159     5.76   0.000     1.047861    1.099671
               |
         _cons |   21.84399   41.73286     1.61   0.106     .5165522    923.7395
--------------------------------------------------------------------------------

. margins, dydx(*)

Average marginal effects                          Number of obs   =        492
Model VCE    : Robust

Expression   : Pr(choice2), predict()
dy/dx w.r.t. : 1.choice1 1.state1 1.state2 1.session2 score

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
   1.choice1 |   .0341974   .0411712     0.83   0.406    -.0464966    .1148914
    1.state1 |  -.0262127   .0494824    -0.53   0.596    -.1231964     .070771
    1.state2 |   .6918165   .0542746    12.75   0.000     .5854404    .7981927
  1.session2 |  -.0112297   .0351127    -0.32   0.749    -.0800493    .0575899
       score |   .0044963   .0010308     4.36   0.000      .002476    .0065166
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

FC TREATMENT

Logistic regression	Number of obs	=	492
	Wald chi2(6)	=	102.34
	Prob > chi2	=	0.0000
Log pseudolikelihood = -108.77105	Pseudo R2	=	0.6757

(Std. Err. adjusted for 41 clusters in subject)

choice2	Odds Ratio	Robust Std. Err.	z	P>|z|	[95% Conf.	Interval]
1.choice1	1.292079	.6226411	0.53	0.595	.5024617	3.322579
1.state1	13.99275	7.383847	5.00	0.000	4.974295	39.36175
1.state2	.0000638	.0001345	-4.58	0.000	1.02e-06	.0039746
1.session2	.5520429	.1616916	-2.03	0.043	.3109271	.9801375
state2#c.score						
0	.9254083	.0185939	-3.86	0.000	.8896732	.9625788
1	1.10163	.0182308	5.85	0.000	1.066471	1.137947
_cons	20.33895	28.34802	2.16	0.031	1.324162	312.4036

```
. margins, dydx(*)

Average marginal effects                          Number of obs   =        492
Model VCE    : Robust

Expression   : Pr(choice2), predict()
dy/dx w.r.t. : 1.choice1 1.state1 1.state2 1.session2 score
```

	dy/dx	Delta-method Std. Err.	z	P>|z|	[95% Conf.	Interval]
1.choice1	.016663	.0320507	0.52	0.603	-.0461552	.0794813
1.state1	.257624	.0581595	4.43	0.000	.1436334	.3716147
1.state2	.6112043	.045856	13.33	0.000	.5213282	.7010805
1.session2	-.0378612	.0193417	-1.96	0.050	-.0757701	.0000478
score	.0010713	.0006895	1.55	0.120	-.00028	.0024226

Note: dy/dx for factor levels is the discrete change from the base level.

NC TREATMENT

```
Logistic regression                               Number of obs   =        492
                                                  Wald chi2(6)    =     195.03
                                                  Prob > chi2     =     0.0000
Log pseudolikelihood = -130.91083                 Pseudo R2       =     0.5968

                                 (Std. Err. adjusted for 41 clusters in subject)
--------------------------------------------------------------------------------
               |               Robust
       choice2 | Odds Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
---------------+----------------------------------------------------------------
     1.choice1 |   6.528634   3.008268     4.07   0.000     2.646093    16.10792
      1.state1 |   1.029548   .4622255     0.06   0.948     .4270629       2.482
      1.state2 |    .000018   .0000266    -7.37   0.000     9.84e-07    .0003276
    1.session2 |    1.38445   .4490417     1.00   0.316     .7331505    2.614338
               |
state2#c.score |
            0  |   .9290356   .0170773    -4.00   0.000     .8961604    .9631167
            1  |   1.120191   .0174057     7.30   0.000     1.086591    1.154831
               |
         _cons |    20.5216   27.46066     2.26   0.024     1.490011    282.6396
--------------------------------------------------------------------------------

. margins, dydx(*)

Average marginal effects                          Number of obs   =        492
Model VCE    : Robust

Expression   : Pr(choice2), predict()
dy/dx w.r.t. : 1.choice1 1.state1 1.state2 1.session2 score

------------------------------------------------------------------------------
             |            Delta-method
             |      dy/dx   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
   1.choice1 |   .1910301   .0509049     3.75   0.000      .0912583    .2908018
    1.state1 |   .0023206   .0358441     0.06   0.948     -.0679325    .0725737
    1.state2 |   .6287987   .0674834     9.32   0.000      .4965337    .7610637
  1.session2 |   .0259479   .0264879     0.98   0.327     -.0259674    .0778633
       score |   .0035593   .0007534     4.72   0.000      .0020827    .0050358
------------------------------------------------------------------------------
Note: dy/dx for factor levels is the discrete change from the base level.
```

T-test of statistical difference between empirical $\Pr_{FC}(\theta_2 = a_1)$ and $\Pr_{NC}(\theta_2 = a_1)$

Two-sample t test with unequal variances

Variable	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf.	Interval]
FC	492	.7682927	.0190411	.4223525	.7308806	.8057048
NC	492	.7398374	.0197993	.4391697	.7009356	.7787392
combined	984	.754065	.0137353	.4308592	.7271112	.7810189
diff		.0284553	.0274696		-.0254506	.0823612

diff = mean(FC) - mean(NC) t = 1.0359

Ho: diff = 0 Satterthwaite's degrees of freedom = 980.507

Ha: diff < 0	Ha: diff != 0	Ha: diff > 0
Pr(T < t) = 0.8497	Pr(|T| > |t|) = 0.3005	Pr(T > t) = 0.1503