Journal of Monetary Economics 000 (2017) 1–16

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Discrete adjustment to a changing environment: Experimental evidence☆

Mel Win Khaw[a], Luminita Stevens[b], Michael Woodford[c,*]

[a] Columbia University, USA
[b] University of Maryland, USA
[c] Department of Economics, Columbia University, 420 W. 118th Street, New York, NY 10027, 212-854-1094, USA

ARTICLE INFO

Article history:
Received 2 December 2016
Revised 4 August 2017
Accepted 4 September 2017
Available online xxx

JEL classification:
C91
D84
E03

Keywords:
Adjustment hazard
Generalized Ss model
Menu costs
Rational inattention

ABSTRACT

A laboratory experiment illustrates cognitive limitations in decision-making that may be relevant for modeling price-setting. Our subjects systematically depart from the optimal Bayesian response in several respects. Their responses are random, even conditioning on available information. Subjects adjust in discrete jumps rather than after each new piece of information, and by both large and small amounts, contrary to the predictions of an "Ss" model of optimal adjustment subject to a fixed cost. And they prefer to report "round numbers," even when that requires additional effort. A model of inattentive adjustment is found to quantitatively outperform popular alternatives.

1. Introduction

A central problem in macroeconomics is understanding the ways in which households and firms respond to changing market conditions, with a particular emphasis on how the immediate (or relatively short-run) effects of new developments differ from the adjustments that eventually occur. When behavior is observed at the micro level, continuous decision variables (such as the price that a firm charges for a given product) are often observed to change only at discrete points in time, even though relevant market conditions are changing continuously; this is often attributed to fixed costs of adjustment ("menu costs" in the case of prices), though there is often little direct evidence about the magnitude of such costs. Here we present evidence for an alternative view, under which such discrete adjustment economizes on the cognitive resources of decision-makers.

This paper attempts to shed light on the nature of discrete adjustment dynamics using a laboratory experiment. While there are obvious questions about the similarities between the task faced by our subjects and those faced in settings of

☆ We would like to thank Fernando Alvarez, Marios Angeletos, Marty Eichenbaum, Eduardo Engel, Benjamin Hebert, Cosmin Ilut, Filip Matějka, Rosemarie Nagel, Erkut Ozbay, Isabel Trevino, Laura Veldkamp, Alex Wolman, and conference and seminar participants at the Barcelona GSE Summer Forum, the NBER Summer Institute, the University of Maryland, the Cowles Summer Conference on Macroeconomics, and the JME-SNB-SCG conference for useful comments, and the National Science Foundation for support of this research.

* Corresponding author.
E-mail address: michael.woodford@columbia.edu (M. Woodford).

http://dx.doi.org/10.1016/j.jmoneco.2017.09.001

relevance for macroeconomic modeling (such as firms' pricing decisions), a laboratory experiment also has important advantages. We can treat the decision-maker's objective as known, to the extent that we assume that our subjects care only about maximizing their monetary payment from the experiment, and we can ensure that many complications sometimes supposed to be relevant for firms' pricing decisions are *not* determinants of our subjects' behavior. We can also be quite certain about exactly what information the decision-maker has at each point in time; not only do we know everything that the experimental subject has had an opportunity to observe, but we know what she ought to understand about the data-generating process, and hence what inferences could rationally be drawn from the information presented. Finally, the laboratory setting allows us to perform multiple replications of the same environment and to vary the environment to draw causal inferences.

Our experiment is a forecasting exercise in which one of two outcomes can occur on each trial, and the subject must estimate the probability of a particular outcome. The underlying probability shifts from time to time, but it is likely to remain the same for many successive trials, offering an opportunity to estimate the current probability from observation of past draws. This kind of exercise allows us to test alternative theories of expectation formation, including the familiar benchmark of "rational expectations." At the same time, the subjects' problem can be viewed as an example of a more general class of problems, in which there is a continuous decision variable (here, the subject's announced probability estimate), with the payoff from a given action depending on the current value of continuous state variable (here, the true probability of an outcome), which varies over time (so that it is necessary to continue monitoring), but is sufficiently persistent to make attempts to monitor the changing state variable worthwhile. Viewed in this way, it is an example of the same basic kind of decision problem faced by a firm that must set a price for its product, where the profits obtained by charging a given price depend on other state variables (demand conditions and factor costs) that fluctuate over time.

In fact, despite the differences in the setting, our experimental data exhibit some notable (and puzzling) features of data on individual prices, as we discuss further below. In particular, our subjects usually leave their decision variable unchanged for periods of time, despite the receipt of many new pieces of information in the meantime, and despite having a continuum of possible responses at each point in time. Since there are no true (external) adjustment costs in our problem, and subjects have (and we believe, are able to understand) all of the information needed to obtain the optimal estimate at each point in time, we conclude that their failure to track the optimal Bayesian estimate more closely reflects imperfect attention, limited memory or related cognitive limitations.[1]

We further demonstrate that our data are consistent with the predictions of a particular quantitative model of inattentive choice in a number of important respects. This model generalizes the model of inattentive discrete adjustment developed in Woodford (2009), most importantly by allowing not only the timing of adjustments but also the choice of where to set the decision variable conditional on adjustment to be inattentive. We compare the quantitative fit of our model with other models of random discrete adjustment, such as the "generalized *Ss* model" of Caballero and Engel (1993, 2007) and the optimizing models with random fixed costs of adjustment proposed by Dotsey et al. (1999) and Dotsey and King (2005). While the latter models describe the adjustment dynamics that we observe better than a simple "*Ss* model," we find that the rational inattention model outperforms these alternatives significantly.

Section 2 describes our experiment. Section 3 gives an overview of some of the salient features of the behavior that we observe, focusing on the ways in which subjects' behavior resembles or differs from the predictions of the ideal Bayesian (or rational expectations) benchmark. Section 4 discusses the extent to which various models of discrete adjustment that have been common in the macroeconomic literature, especially the literature on price adjustment, can account for our data. Section 5 offers a concluding discussion.

2. The experimental design

Our laboratory experiment follows the setup of Gallistel et al. (2014), who study how well subjects can predict probabilities.[2] The subjects' task is to estimate the probability that the next draw is a green ring from a box with both green and red rings. The subjects draw rings with replacement from the box and indicate their draw-by-draw estimate. Fig. 1 shows a screenshot of the experiment. The screen displays the box with a hidden number of red and green rings on the left. The slider at the bottom indicates the subject's estimate on the current trial, $\hat{p}_t$, with the number in percentage points displayed above the bar. The box on the right side of the screen displays 1000 rings that also depict the subject's estimate visually. Whenever the subject moves the slider, this box is adjusted in real time to reflect the probability indicated by the slider position. The subject begins with a guess and adjusts the position of the slider to indicate his or her estimate. Each time the subject clicks the NEXT button, a new ring is drawn randomly from the box and displayed on the screen. The subject may then adjust his or her forecast or leave it unchanged. After each ring draw, the subject's cumulative score is updated and displayed at the top of the screen. This process is repeated until the session ends, after 1000 ring draws.

[1] Our experimental design differs importantly from that of Magnani et al. (2016), who study how well an "*Ss*" model fits laboratory data on the timing of adjustments. Their experimental setup imposes an external fixed cost of adjustment, in order to ensure that adjustment will be discrete, but considers whether adjustments are optimally timed; we are instead interested in observing discrete adjustment even when subjects are free to adjust at any time. Their setup also requires that when adjustment occurs, the decision variable is moved to exactly the currently optimal state, whereas we are interested in where the decision variable will be set when subjects are free to set it anywhere.

[2] We modify their experiment in a number of respects, because of the different focus of our study, as discussed in the online appendix.

Please cite this article as: M.W. Khaw et al., Discrete adjustment to a changing environment: Experimental evidence, Journal of Monetary Economics (2017), http://dx.doi.org/10.1016/j.jmoneco.2017.09.001