European Economic Review 98 (2017) 1–31

Contents lists available at ScienceDirect

European Economic Review

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

Response functions☆

Carlos Oyarzun[a,*], Adam Sanjurjo[b], Hien Nguyen[c]

[a] School of Economics, The University of Queensland, Brisbane, Queensland 4072, Australia
[b] Department Fundamentos del Análisis Económico, Universidad de Alicante, Spain
[c] Department of Mathematics and Statistics, School of Engineering and Mathematical Sciences, La Trobe University, Australia

ARTICLE INFO

Article history:
Received 31 December 2015
Accepted 11 June 2017
Available online 20 June 2017

JEL Classification:
D81
D83

Keywords:
Adaptive learning
Response functions
Experimental economics
Stochastic dominance

ABSTRACT

Imagine that John must choose between two uncertain payoff distributions, knowing that the set of possible payoffs is the same for both, but nothing about the shapes of the distributions. In the first period he chooses either alternative and experiences a payoff as a result of his choice. Given this experienced payoff, in the second period he decides whether to choose the same alternative again, or switch. We model John's second period choice with a *response function*, i.e., a mapping from obtained payoffs to the probability of choosing the same alternative in the second period. We first provide results on (i) how the shape of the response function affects both expected payoffs and exposure to risk, and (ii) what standard models of choice under uncertainty would predict about the shape of the response function. We then run an experiment to elicit subjects' response functions, empirically characterize the heterogeneity across subjects with a mixture model, and illustrate how payoffs vary across response function types. Finally, we use our theoretical results, along with additional information that we collected from subjects, to interpret their response functions.

1. Introduction

Imagine that John walks up to two slot machines. By dropping a coin in either machine he will receive a prize of between one and five gold bars. However, he knows nothing about the probability distribution over prizes for either machine. John has exactly two coins in his pocket. Upon choosing a machine for the first coin, and experiencing the corresponding prize, he must now decide whether to drop his second coin in the same machine, or switch. While simple, this environment

☆ This paper has benefited from useful discussions with Carlos Alós-Ferrer, Vincent Crawford, Carlos Cueva, Nobuyuki Hanaki, Jonas Hedlund, Ed Hopkins, Nagore Iriberri, Kenan Kalaycı, Vadym Lepetyuk, Rosario Macera, Geoff McLachlan, Friederike Mengel, Joshua Miller, Andreas Ortmann, Marciano Siniscalchi, and audiences at the 2016 ANZWEE, 2015 Econometric Society World Congress, 2015 ESA World Meetings, 2015 Australian Conference of Economists, SABE 2014 Lake Tahoe, Pontificia Universidad Católica de Chile, Universidad de Alicante, Universidad de Concepción, University of New South Wales, and University of Technology Sydney. We thank three referees and an Associate Editor for insightful and constructive comments. We also acknowledge financial support from the Spanish Ministry of Sciences and Technology (SEJ2007-62656), the Spanish Ministry of Economics and Competition (ECO2012-34928), and the University of Queensland Start-up Grant. All errors remain our responsibility.

* Corresponding author.
E-mail address: c.oyarzun@uq.edu.au (C. Oyarzun).

http://dx.doi.org/10.1016/j.euroecorev.2017.06.011
0014-2921/

approximates important economic decisions in which virtually no information is known ex-ante about the shape of payoff distributions, so direct experience (while limited) may have a large effect on future choice.[1,2]

Since the first period choice is made with no information, our analysis treats it as exogenously given. On the other hand, the second period choice may depend on the payoff experienced as a result of the first period choice. Thus, we represent the second period choice as a Bernoulli distribution in which the probability of staying with the same alternative is determined by the first period payoff. This representation defines a *response function*, i.e., a function mapping the experienced payoff to the probability of choosing the same alternative in the second period.

In principle, the response function is a belief-free model of choice, similar to those found in adaptive-learning models of choice, such as Easley and Rustichini (1999), Börgers et al. (2004), Mengel and Rivas (2012), Oyarzun and Sarin (2013), and Agastya and Slinko (2015). It can be viewed as the “primitive” reaction of an individual to obtained payoffs, in a problem in which decisions are made not “from description” of how likely the payoff consequences of different alternatives are, but “from experience” (cf. Erev and Haruvy (2013)). An alternative interpretation that we will also consider is that the response function is a consequence of a more structured (belief-based) decision approach such as Bayesian Expected Utility Theory.[3]

Regardless of whether one interprets the response function as a belief-free model of choice, or as having an underlying Bayesian structure, the following fundamental questions arise in the setting we consider: when people are so poorly informed, and have so little opportunity to learn, does it even matter how they respond? In particular, are there general patterns in the way that one particular response function or another affects payoffs? It turns out that there are. We first analyze: (i) how the shape of the response function affects both expected payoffs and exposure to risk, and (ii) what standard models of choice under uncertainty would predict about the shape of the response function. We then run an experiment to elicit subjects’ response functions, empirically characterize the heterogeneity across subjects with a mixture model, and illustrate how payoffs vary across response function types. Finally, we use our theoretical results, along with additional information that we collected from subjects, to interpret their response functions.[4]

We begin our theoretical analysis in Section 2.1 by assuming that payoff distributions can be ordered according to stochastic dominance. Proposition 1 establishes the intuitive result that if an individual’s response function is increasing (concave), then the alternative she chooses in the second period is more likely to first-order (second-order) stochastically dominate the unchosen alternative.[5] Then, we drop the assumption that distributions are ordered according to stochastic dominance, and provide an analogous result. In particular, Proposition 3 shows that if an EU-maximizer’s response function is an affine transformation of her Bernoulli utility function, then the alternative she chooses in the second period is more likely to be the alternative that she would prefer if both payoff distributions were known.

Given the benefits that can arise from the increasingness of an individual’s response function, one can ask whether stretching response functions to increase their slope can result in further improvements in economic performance. Accordingly, Proposition 2 shows that if an individual’s response function is of *maximum strength*, meaning that it stays (switches) with probability one if the obtained payoff is the highest (lowest), then no other response function can outperform it in the second period for all pairs of payoff distributions. Further, Corollary 1 shows that no response function outperforms all others for all pairs of payoff distributions. Finally, Corollary 2 shows that an EU-maximizer whose response function is an affine transformation of her Bernoulli utility function is expected to draw from her preferred payoff distribution most often when the affine transformation is stretched to maximum strength.

[1] A few examples are the choice of realtor when selling a house, of a lawyer for litigation, or of a school for each child. Similarly, investment decisions, such as whether to participate in the IPO of a certain company and subsequent secondary offerings, may not be available often, giving investors only limited opportunities to learn from their own experiences.

[2] Experimental evidence indicates that decision-makers may often have a clearer understanding of what outcomes may occur than the probability distributions over those outcomes (see, e.g., Loomes (1998), Selten et al. (1999)), and systematically over-weight experienced relative to observed information (see, e.g., Maniadis and Miller (2012), Simonsohn et al. (2008)).

[3] Under this interpretation, in the second period the individual chooses the alternative that yields the highest expected utility according to her updated beliefs.

[4] A natural question is to what extent choice, in the setting that we consider, can be rationalized in terms of axioms on the response function. While, in principle, this exercise would be helpful for developing descriptively-better models of preferences over ambiguous alternatives, the extreme minimality of information in our setting rules out the possibility of verifying the axioms that arise in relatively richer information environments (e.g., Anscombe and Aumann (1963), Gilboa and Schmeidler (1989), Klibanoff et al. (2005)). For instance, Anscombe and Aumann (1963) assume that when individuals choose from ambiguous lotteries that differ only in the prize associated with one state of the world, then preferences over lotteries are fully determined by preferences over prizes in that state. By contrast, because in our setting individuals lack the type of information available in theirs (i.e., how acts map states to prizes), we are unable to provide an analogous axiom. Similarly, the belief-free approaches referenced earlier tend to also rely on a relatively greater availability of information. For example, there it is typical to consider an individual who acquires information over time, by repeated experience, which is not possible in the one-shot learning environment that we consider. To illustrate, Easley and Rustichini (1999) consider axioms regarding how preferences over alternatives change in response to observed states, and assume “exchangeability,” i.e., that upon observing two states, preferences over alternatives should be the same, regardless of what order the states are observed in. An analogous axiom is not possible in our setting because the individual observes only one payoff realization. On the other hand, the analyses of Börgers et al. (2004) and Oyarzun and Sarin (2013), as explained below, relate more closely to our approach.

[5] It has long been observed in the psychology and computer sciences literatures that learning models which increase the probability of choosing those alternatives that have been successful, have the property that the “ex ante” expected value of the probability of choosing the alternative that is most likely to result in a success is increasing in time (see, e.g., Norman (1968) and Narendra and Thathachar (1974)). These insights were first introduced to economics by Börgers et al. (2004), and Oyarzun and Sarin (2013) extended their analysis to include risk aversion.