Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 162 (2016) 352–371

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Belief revision generalized: A joint characterization of Bayes' and Jeffrey's rules ☆

Franz Dietrich [a,*], Christian List [b,*], Richard Bradley [b,*]

[a] *Paris School of Economics & CNRS, France*
[b] *London School of Economics, United Kingdom*

Received 1 August 2013; final version received 6 August 2015; accepted 12 November 2015

Available online 2 December 2015

Abstract

We present a general framework for representing belief-revision rules and use it to characterize Bayes' rule as a classical example and Jeffrey's rule as a non-classical one. In Jeffrey's rule, the input to a belief revision is not simply the information that some event has occurred, as in Bayes' rule, but a new assignment of probabilities to some events. Despite their differences, Bayes' and Jeffrey's rules can be characterized in terms of the same axioms: *responsiveness*, which requires that revised beliefs incorporate what has been learnt, and *conservativeness*, which requires that beliefs on which the learnt input is 'silent' do not change. To illustrate the use of non-Bayesian belief revision in economic theory, we sketch a simple decision-theoretic application.

JEL classification: C73; D01; D80; D81; D83; D90

Keywords: Belief revision; Subjective probability; Bayes' and Jeffrey's rules; Axiomatic foundations; Fine-grained versus coarse-grained beliefs; Unawareness

☆ We are grateful for comments from the editors and referees and from audiences at D-TEA 2010 (HEC & Polytechnique, Paris, June 2010), LSE's Choice Group Seminar (LSE, September 2010), Pluralism in the Foundations of Statistics (University of Kent, September 2010), and Decisions, Games, and Logic 2012 (University of Munich, June 2012). Although this paper is jointly authored, List and Bradley wish to note that the bulk of the mathematical credit should go to Dietrich. Dietrich was supported by a Ludwig Lachmann Fellowship at the LSE and the French Agence Nationale de la Recherche (ANR-12-INEG-0006-01). List was supported by a Leverhulme Major Research Fellowship (MRF-2012-100) and the Franco-Swedish Program in Philosophy and Economics (via a visit to the University of Uppsala). Bradley was supported by the Arts and Humanities Research Council (via a grant on 'Managing Severe Uncertainty', AH/J006033/1).

* Corresponding authors.
E-mail addresses: fd@franzdietrich.net (F. Dietrich), c.list@lse.ac.uk (C. List), r.bradley@lse.ac.uk (R. Bradley).

http://dx.doi.org/10.1016/j.jet.2015.11.006

1. Introduction

A belief-revision rule captures how an agent's subjective probabilities should change when the agent learns something new. The standard example is Bayes' rule. Here, the agent learns that some event has occurred, and the response is to raise the subjective probability of that event to 1, while retaining all probabilities conditional on it. Formally, let Ω be the underlying set of possible worlds (where Ω is non-empty and, for simplicity, finite or countably infinite).[1] Subsets of Ω are called *events*. Beliefs are represented by some probability measure on the set of all events. Bayes' rule says that, upon learning that some event $B \subseteq \Omega$ has occurred (with $p(B) \neq 0$), one should move from the prior probability measure p to the posterior probability measure p' given by

$$p'(A) = p(A|B) \text{ for all events } A \subseteq \Omega.$$

In economic theory, belief changes are almost always modelled in this way. The aim of this paper is to draw attention to a more general form of belief revision, which is seldom discussed in economics. We develop a general framework in which different belief-revision rules – Bayesian and non-Bayesian – can be characterized. In this framework, the key difference between different belief-revision rules lies in what they take to be the input prompting the agent's belief change. Under Bayes' rule, the learnt input is always the occurrence of some event, but this is more restrictive than often recognized, and we show that there is scope for useful generalization.

We begin with an example that will resonate with any international traveller. Kotaro, a junior academic from Japan, is applying for a faculty position in the UK. After his interview, he is telephoned by Terence, the chair of the department, to inform him of the outcome. Kotaro gets a vague impression that he is being offered the job, but struggles to understand Terence's thick Irish accent. At the end of the call, he is still unsure whether he has received an offer. He becomes convinced only after a subsequent conversation with another member of the department.

This example illustrates an instance of belief revision triggered by a noisy signal. Before the telephone conversation with the chair of the department, Kotaro attaches a very low probability to the event of getting the job. After the conversation, he attaches a somewhat higher probability to it, but one that still falls short of certainty. For this, a second conversation (with another person) is needed. Such cases present challenges to the Bayesian modeller. If the first change in Kotaro's probabilities is to be modelled as an application of Bayes' rule, then it will clearly not suffice to restrict attention to the 'naïve' set of possible worlds $\Omega = \{\text{appointed, not appointed}\}$. Relative to that 'naïve' set, a Bayesian belief change could never increase the probability of the event 'appointed' without raising it all the way to 1.

The modeller will need to enrich the set Ω to capture the possible sensory experiences responsible for Kotaro's shift in probabilities over the events 'appointed' and 'not appointed'. So, the enriched set of worlds will have to be something like

$$\Omega' = \{\text{appointed, not appointed}\} \times \mathcal{A},$$

where $\mathcal{A}$ is the set of possible analogue auditory signals received by Kotaro's eardrums. The signal he receives from Terence will then correspond to some subset of Ω', specifically one of the form $B = \{\text{appointed, not appointed}\} \times A$, where $A \subseteq \mathcal{A}$ is a particular 'auditory event'.

[1] We expect that our results can be generalized to a set Ω of arbitrary cardinality.