WORKING PAPERS

N° TSE-574

May 2015

"Dinner with Bayes: On the Revision of Risk Beliefs"

Christoph M. Rheinberger and James K. Hammitt

Toulouse School of Economics

Dinner with Bayes: On the Revision of Risk Beliefs

Christoph M. Rheinberger and James K. Hammitt*

March 9, 2015

Abstract

We study how people form and revise risk beliefs based on food safety information. In an online experiment, subjects stated their perceived risk of contracting a foodborne illness before and after receiving information about the eating habits of the average consumer. The majority of subjects revised their risk beliefs consistent with the Bayesian learning hypothesis. Personal consumption habits affected belief revisions in the expected direction, while precautionary behavior entered the updating process through the priors. Fewer than 20% of the subjects responded inconsistently to the information provided. Several factors related to the subject's numerical skills explain inconsistent belief revisions. *JEL I12, I18, D80*

*CMR: Toulouse School of Economics (LERNA-INRA), email: rheinberger.cm@gmail.com. JKH: Harvard Center for Risk Analysis and Toulouse School of Economics (LERNA-INRA), email: jkh@harvard.edu. The research in this paper was sponsored by an INRA Scientific Package and a grant of the RHW Foundation, Liechtenstein. JKH acknowledges additional support from the European Research Council (FP7/2007-2013, grant no. 230589). We thank Henrik Andersson and participants of the Risk, Perception, and Response Conference at the Harvard Center for Risk Analysis for useful comments. All remaining errors are ours.

1 Introduction

People often respond to public health policies in ways that are inconsistent with economic theory. They overreact to some risks while they ignore others (Slovic et al. 2000); they are reluctant to give up unhealthy behaviors though they know it would be better for them (O'Donoghue and Rabin 2001); and they take healthy behaviors as excuse for indulging in unhealthy behaviors. Examples include eating more when foods are low in calories (Wisdom et al. 2010) and smoking cigarettes down to the bone while cutting back on cigarettes (Adda and Cornaglia 2006).

Reasons for such obvious deviations from the rational consumer model are manifold (McFadden 2006) and include cognitive and attentional limitations, emotional arousal, various forms of procrastination, and difficulties in processing health information (Cawley and Ruhm 2011). In this paper, we focus on the processing of information about health risks. In particular, we study how consumers perceive the risk of contracting a foodborne illness before and after the provision of risk-related information. Risk perception is a critical link in the causal chain between consumer information and behavioral responses (Magat and Viscusi 1992; Viscusi 1998; Sloan et al. 2003). A better understanding of how consumers form perceptions of food risks and how they adjust their beliefs to new information is therefore of considerable interest to policymakers. The interest is fueled by its implications for the evaluation of existing food safety policies and by regulatory needs to accurately predict behavioral responses to new consumer information and awareness campaigns. Success or failure of such campaigns matters because food pathogens cause billions of episodes of foodborne illness each year and the corresponding welfare costs are tremendous.[1]

Two questions emerge. Do public information programs affect health risk perceptions? And if so, do they alter consumer behavior in the predicted manner? Again, answers to these questions require a better understanding of

[1] In the U.S. alone, the annual social cost of foodborne illnesses is more than $50 billion (Scharff 2012).

the processing of risk-related information. The uptake of information plays a key role in studies of risk perception (Slovic et al. 2000). Both economists (e.g., Arrow 1982; Viscusi 1997; McFadden 1999) and psychologists (e.g., Johnson and Tversky 1983; Loewenstein et al. 2001; Kahneman 2003; Slovic et al. 2004) have long recognized that people make a number of common mistakes when they update risk beliefs with newly available information: small risks are overestimated while large ones are underestimated, risks are assessed based on emotions rather than cognitive evaluations, and more attention is given to bad news than to good news.[2]

Over the past 30 years, a number of empirical studies have addressed the impact of information on subjective risk beliefs. Viscusi and O'Connor (1984) elicited chemical workers' perception of job hazards based on warning labels. Smith and Johnson (1998) studied the effect of public information programs on homeowners' attitudes toward the health risk associated with radon exposure. Dickie and Gerking (1996) explored public perceptions of skin cancer. Viscusi (1997) studied location decisions in the presence of ambiguous information about air pollution. More recently, Cameron (2005) tested how conflicting information affects the perceived risk of climate change; Lundborg and Lindgren (2004) and Viscusi and Hakes (2008) examined beliefs about lung cancer risk and smoking behavior in Sweden and the U.S., respectively; Andersson and Lundborg (2007) compared the perception of road-traffic and overall mortality risk in Sweden.

Most individuals in the above studies revised their risk beliefs in the expected direction. However, the studies largely ignored the endogenous nature of health risks. Indeed, people often have private information about their health and take precautionary measures to reduce the likelihood or severity of bad health outcomes (Shogren and Stamland 2007). Such measures are typically unobserved in observational studies, but may systematically affect both perceptions and actions. For example, consumers choose the quality, storage

[2]As people process information in accordance with their feelings about a particular risk (Loewenstein et al. 2001; Slovic et al. 2004), it is of little surprise that negative news attract more attention. The emotional filtering is apparently reversed when subjects form and revise beliefs about desirable aspects of life (Eil and Rao 2011; Mobius et al. 2011).

place, and preparation of their foods, thereby affecting the likelihood of contracting a foodborne illness. It might thus be perfectly rational for a consumer to hold a subjective risk belief that differs from the statistical population risk.

This paper presents a belief-elicitation protocol that permits capturing the impact of precautionary behaviors and other idiosyncratic characteristics affecting both the formation and the revision of subjective risk beliefs. In what is essentially a panel structure, a representative sample of French consumers stated their perceived chance of contracting a foodborne illness from eating fish. We first elicited subjects' risk beliefs without any specific information. We repeated the elicitation after having provided subjects with information about the average consumer's fish consumption, the corresponding population average risk, and various risky and risk-averting behaviors. The chained elicitation procedure allows us to explore subjects' responses to risk-related information, heterogeneity in the revision of risk beliefs, and deviations from the Bayesian rationality assumption that underlies the design of most—if not all—consumer information campaigns.

In a nutshell, we find that the majority of subjects updated their beliefs consistently with the Bayesian learning hypothesis. These subjects responded to information about the statistical risk by reducing their prior risk beliefs if these were above the statistical risk and by increasing their prior risk beliefs if they were below the statistical risk. This finding holds whether or not we control for behavioral and/or socioeconomic factors. Precautionary effort in handling and preparing food had the expected attenuating effect on the formation of prior risk beliefs, but did not directly affect the updating process. This result underpins the importance of controlling for confounding factors in the formation of risk beliefs and has several implications for predicting the outreach of existing health and consumption advisories.

Using finite mixture models, we decompose the heterogeneity in belief revisions and find four distinct updating patterns: (1) subjects who ignore the new information altogether; (2) subjects who take the information at face value (i.e. accept the statistical risk as theirs); (3) subjects who combine the provided information and their prior to update beliefs; and (4) subjects who update in an inconsistent manner. The mixture modeling approach allows us to link the emerging patterns back to personal characteristics. In particular, we find that older, less educated and less literate subjects are more likely to adapt either strategy (1) or (4) when updating their risk beliefs. As we shall discuss in more detail both refusal of information and lack of numeracy are problematic from the regulator's point of view as they undermine the efficacy of public health policies that seek to change behavior by informing consumers.

The paper proceeds as follows. In Section 2, we operationalize the Bayesian learning model and derive a formal definition of rational updating that is conditional on the precautionary effort expended by the updater. Section 3 provides details of the belief-elicitation task and the sample characteristics. Section 4 outlines our econometric approach. Section 5 presents the results of our study. In Section 6, we discuss the response to information both at the individual and the aggregate level. Section 7 concludes.

2 Bayesian Learning Model

Ample evidence from both experimental and observational studies suggests that people overestimate the likelihood of rare events and underestimate the likelihood of frequent ones (Barron and Yechiam 2009). Since the seminal paper by Lichtenstein et al. (1978) dozens of studies have shown that this observation specifically applies to the context of health risks, with people being either overly optimistic or pessimistic about their risk of dying or of developing a specific disease.[3]

Economic theory holds that accurate information about the nature of the risk and the means of precaution may help people to better align their beliefs to the actual risk. Yet in the real world people might—willingly or unwillingly—ignore information. Viscusi (1989) assumed that subjects do not treat the probabilities presented to them as fully informative and proposed a model in which individuals use probabilistic information in a Bayesian fashion

[3]Harris and Hahn (2011) warn that many of these studies might be methodologically flawed.

to revise their risk beliefs.[4] He argued that the Bayesian updating process is consistent with two possible interpretations: individuals might not have full confidence in the source of information, or they might simply treat any risk-related information as not perfectly applicable to their individual circumstances.

2.1 Theoretical Model

Both interpretations allow people to discount the new information within the updating process. This can be formalized in the most basic version of the Bayesian learning model:

$$q_i = \frac{\gamma p_i + \xi s}{\gamma + \xi} = \gamma^* p_i + \xi^* s, \tag{1}$$

where q_i denotes individual i's posterior risk belief, which is formed based on i's prior risk belief p_i and the information about the statistical risk s; γ and ξ are the information contents associated with p_i and s, respectively; and $\gamma^* = \gamma/(\gamma+\xi)$ and $\xi^* = \xi/(\gamma+\xi)$ are the corresponding precision weights.

Equation (1) assumes that individuals form their posterior belief as a weighted average of the belief they held prior to receiving the risk-related information and the inference drawn from the new information. A limitation of the basic Bayesian learning model is that it treats the interpretation of new information as a black box. Smith and Johnson (1988) proposed a generalization of the basic model in which factors that affect the inference process might also influence people's perception of the relative precision of either their prior beliefs or the information content. As Smith and Johnson note, it is likely that some of these factors also affect the formation of prior risk beliefs.

Based on these insights, the Bayesian learning model can be extended to explore heterogeneity in the response to risk-related information. We assume that people form their posterior risk belief by processing new pieces of information and combining them with knowledge of personal exposure and

[4] Similar models have been used by psychologists since the early 1950s (see Slovic and Lichtenstein 1971).

precautionary behavior. This leads to the behavioral model:

$$q_i = \gamma^*(\mathbf{A}_i, \theta_\mathbf{A}) p(\mathbf{B}_i, \theta_\mathbf{B}) + \xi^*(\mathbf{C}_i, \theta_\mathbf{C}) s(\Delta\mathbf{D}_i, \theta_\mathbf{D}), \tag{2}$$

where $\theta_\bullet$ are parameter vectors. The weighting functions $\gamma^*(\mathbf{A}_i, \theta_\mathbf{A})$ and $\xi^*(\mathbf{C}_i, \theta_\mathbf{C})$ are contingent upon factors (summarized in vectors $\mathbf{A}_i$ and $\mathbf{C}_i$) that influence individual perception of the relative precision of the prior and the information, respectively.

The prior risk belief $p(\mathbf{B}_i, \theta_\mathbf{B}) \equiv p_i$ is a function of personal factors (age, gender, education, etc.) collected in the vector $\mathbf{B}_i$. Similarly, personal inference from the received risk information $s(\Delta\mathbf{D}_i, \theta_\mathbf{D}) \equiv s_i$ is a function of behavioral factors (exposure, precautionary effort, etc.). Instead of directly including the latter factors, we measure individual i's behavioral deviation from the average consumer by the vector $\Delta\mathbf{D}_i$. This deviation is crucial for individual i's interpretation of the statistical risk s. If i believes herself to be more (or less) exposed to a specific risk than the average person, she uses s as a reference point for adjusting her prior belief (Hogarth and Einhorn 1992).

We invoke the behavioral implications of equation (2) to define what is a rational response to risk-related information conditional on the consumer's behavior. Consider individual i's response to information about the average consumer's behavior and about the corresponding statistical risk s. Given the consumer's observed prior risk belief p_i, we define consistent updating by the following set of behavioral rules.

Definition. *A consistent response to risk-related information does not violate any of the following conditions: (i) if $p_i = s$ and $\Delta\mathbf{D}_i = 0$, then $s = p_i = q_i$; (ii) if $p_i > s$ and $\Delta\mathbf{D}_i = 0$, then $s \leq q_i \leq p_i$; (iii) if $p_i < s$ and $\Delta\mathbf{D}_i = 0$, then $s \geq q_i \geq p_i$; (iv) if $p_i \geq s$ and $\Delta\mathbf{D}_i > 0$, then $s \leq q_i$; (v) if $p_i \leq s$ and $\Delta\mathbf{D}_i < 0$, then $s \geq q_i$; (vi) if $p_i > s$ and $\Delta\mathbf{D}_i < 0$, then $p_i \geq q_i$; (vii) if $p_i < s$ and $\Delta\mathbf{D}_i > 0$, then $p_i \leq q_i$.*

Conditions (i)-(iii) provide a behavioral reformulation of equation (1) that applies to individuals who behave like the average consumer; conditions (iv)-(vii) prescribe how an individual should respond when they engage in more or less risky behavior than the average consumer. Figure 1 maps out all possible belief revisions. Compliance with the first three conditions corresponds to subjective belief revisions in the gray-shaded area. Compliance with the latter four conditions suggest that belief revisions outside the gray-shaded area are rational only if the subject differs from the average consumer in a prescribed manner. Hence, the definition classifies any dynamically consistent belief revision as rational (Savage 1954). In the empirical part, we will make use of the distinction between consistent and inconsistent responses to risk-related information.

Figure 1: Landscape of possible belief revisions

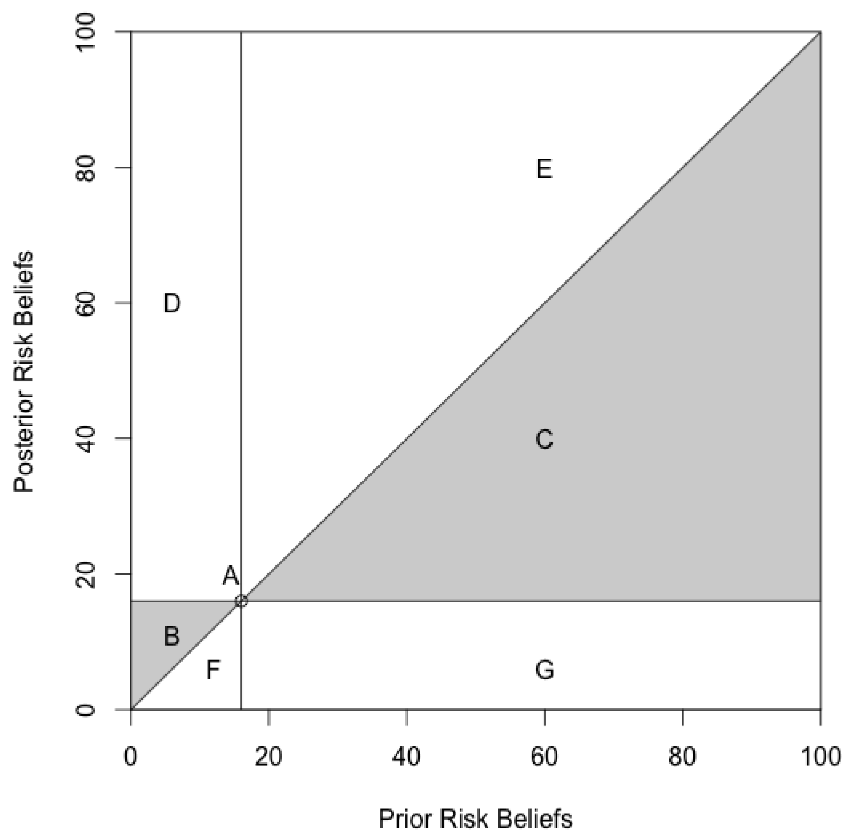

Notes: Each of the conditions of the consistency definition delimits a specific area: condition (i) ~ A; condition (ii) ~ C; condition (iii) ~ B; condition (iv) ~ C ∪ E; condition (v) ~ C ∪ G; condition (vi) ~ B ∪ F; condition (vii) ~ B ∪ D.

2.2 Empirical Model

In practice, we often do not have sufficient information about the detailed behavioral processes captured by equation (2) to estimate the parameters of interest. Yet under reasonable assumptions, an empirically testable version of the extended Bayesian learning model can be obtained (Smith and Johnson 1988). First, we presume substantial overlap between $\mathbf{A}_i$, $\mathbf{B}_i$, and $\mathbf{C}_i$, and stack them into the single vector $\mathbf{X}_i$. Second, we impose a linearly additive form for each of the behavioral functions in equation (2). Upon appending a stochastic error term, we obtain a workhorse model to analyze observed heterogeneity in updating behavior:

$$q_i = \beta_0 s + \beta_1 p_i + \beta_2' \Delta\mathbf{D}_i + \beta_3' \mathbf{X}_i + \beta_4' p_i \Delta\mathbf{D}_i + \beta_5' p_i \mathbf{X}_i + \beta_6' \Delta\mathbf{D}_i \mathbf{X}_i + \beta_7' p_i \Delta\mathbf{D}_i \mathbf{X}_i + \epsilon_i. \quad (3)$$

Equation (3) allows us to capture the possibility that some factors which affect the belief revision process may also affect the formation of the prior risk belief.

3 Experimental Design

Two premises guided the development of our belief-elicitation protocol: (1) People are not very good in making sense of small probabilities, but (2) they do fairly well in reporting expectations for specific states of the world as a percent chance (Manski 2004). The belief-elicitation task proceeded as follows. We first informed subjects about the annual number of cases of foodborne illness in France (about 250,000). With this information in hand, they stated on a semi-quantitative scale how frequently they expected to suffer a foodborne illness.[5] That is they gave us an estimate of their personal risk of contracting a foodborne illness. Next, we instructed subjects to assume they will suffer a foodborne illness sometime next year, and inquired how likely they thought it was (in terms of a percent chance) that the cause for their illness would be

[5]The scale included eight ordered categories, ranging from once per month to less than once in a lifetime.

from eating fish.[6] In other words, we elicited the conditional risk of a fishborne illness. The task was computer-based and subjects indicated their conditional risk estimates using a percent slider.

In a following step, subjects received information about the fraction of cases of foodborne illness in France attributable to eating fish, the consumption habits of French fish consumers, and behaviors that may reduce or increase the risk of contracting a foodborne illness.[7] We asked them to consider this information when revising their prior risk estimate. This time the percent slider had additional marks indicating the subject's prior and the statistical risk. Based on the information provided to them, subjects who knew the approximate population of France (about 66 million) could calculate the population annual average risk of foodborne illness ($250,000/66$ million $\approx 4/1,000$) and the risk of illness from eating fish ($16\% \times 250,000/66$ million $\approx 6/10,000$). Moreover, we can infer the subject's belief about her absolute risk of illness from eating fish based on the elicited risk beliefs. Any deviation from the population average conditional risk can hence be explained by her specific patterns of preparation and consumption of fish or other foods and her beliefs about the absolute risk of the types of fish and other foods she consumes.[8] While we

[6] Following Manski (2004, p.1343) we reminded subjects that "*the percent chance must be a number from 0 to 100. Numbers like 2 or 5 percent may be 'almost no chance', 20 percent or so may mean 'not much chance', a 45 or 55 percent chance may be a 'pretty even chance', 80 percent or so may mean a 'very good chance', and a 95 or 98 percent chance may be 'almost certain'.*"

[7] Based on epidemiological data for France (Vaillant et al. 2005) we presented the conditional risk that a foodborne illness is from eating fish as $s = 16\%$.

[8] Eq. (1) can be re-written so that the conditional risk of a fishborne illness is broken up into the absolute risk of contracting a fishborne illness F and the absolute risk of contracting a foodborne illness N:

$$q_i = \gamma^* p_i + \xi^* s \leftrightarrow \frac{F_q}{N_q} = \gamma^* \frac{F_p}{N_p} + \xi^* \frac{\bar{F}}{\bar{N}}.$$

The experimental procedure consisted of the following steps: (1) subjects were told $\bar{N}$; (2) conditional on $\bar{N}$ they formed an estimate of N_p; (3) they provided an estimate of $p = \frac{F_p}{N_p}$; (4) they received information about the statistical conditional risk $s = \frac{\bar{F}}{\bar{N}}$; (5) based on $p = \frac{F_p}{N_p}$ and $s = \frac{\bar{F}}{\bar{N}}$ they provided an estimate of $q = \frac{F_q}{N_q}$. For a Bayesian belief updater $N_q = N_p$, unless he based his assessment of $N_q = N_p$ on false assumptions about the average consumer's fish consumption and/or precautionary effort. Consider the following example.

enquired about fish consumption and preparation methods, our information about these factors is necessarily limited.

The belief-elicitation task was included in a large online survey and administered to a French consumer panel during July and September 2012. We obtained answers from 1,003 panel members who eat fish at least once a week. As the sample matches quotas for age, gender, region, and employment status, we take it to be representative of French consumers.[9] Apart from the usual socioeconomic indicators we inquired about the quantity of consumed fish; the preferred purchase, storage and preparation methods; the importance of various self-protection behaviors; and other attributes that may determine food consumption choices (Shogren and Stamland 2007). Table 1 summarizes our data.

Subject i beliefs his risk is $p = 10\% = \frac{F_p}{N_p} = \frac{10{,}000}{250{,}000}$, but he formed his belief about N_p based on the assumption that the average consumer eats fish once a week. Being told that the average consumer eats fish three times a week may let i to revise his belief about the risk of contracting a foodborne illness to $N_q \geq N_p$, because the information provided links the average absolute risk to behavioral factors. Our definition of consistent updating accounts for such inferences drawn from the average consumer's behavior, only requiring belief revisions to be dynamically consistent in a Savagean sense (e.g., $N_q < N_p$ would be an inconsistent inference from the information provided).

[9]Quotas were based on 2009 French census data (http://www.insee.fr/fr/recensement-2009.htm). The required frequency of fish consumption imposed an additional eligibility constraint resulting in a response rate of roughly 25%.

Table 1: Sample characteristics

Variable		Obs.	Mean	St. Dev.	Min.	Max.
PRIOR	Prior risk belief (p)	1,003	0.32	0.25	0	0.99
POSTERIOR	Posterior risk belief (q)	1,003	0.23	0.20	0	0.93
BASELINE PESSIMISM	Prior risk belief larger than statistical risk ($p > s$)	1,003	0.48	0.49	0	1
CONTROL PESSIMISM	Posterior risk belief larger than statistical risk ($q > s$)	1,003	0.48	0.50	0	1
MALE	Gender is male	1,003	0.49	0.50	0	1
KIDS	Lives with children younger than 10 years	1,003	0.28	0.45	0	1
AGE	Subject's age	1,003	43.56	13.48	18	80
INCOME	Monthly household income in €1,000	950	2.84	0.12	0.25	5
EDUCATION	Years of education	1,003	14.20	3.41	5	17
PREGNANCY	Pregnancy within the subject's household	1,003	0.10	0.31	0	1
HEALTH STATUS	Self-assessed health status[a]	1,003	7.36	1.71	0	10
SAFETY CONCERNS	Subject has concerns over the safety of seafood	1,003	0.48	0.5	0	1
NUMERACY	Subject passed a simple numeracy test[b]	1,003	0.93	0.25	0	1
SUBJECTIVE RISK	Subjective risk of contracting a foodborne illness per life-year	1,003	0.41	0.45	0	1
RAW FISH	Number of times subject consumes raw fish[a]	1,003	0.95	1.73	0	10
WASH HANDS	Number of times hands are washed before meal preparation[a]	1,003	7.20	3.87	0	10
STORE FISH	Number of times fish is kept $\geq$ 3 days in the fridge[a]	1,003	0.83	2.14	0	10
PREPARE FISH	Number of times subject prepares fish well done [a]	1,003	5.96	4.07	0	10
MEALS	Number of times subject eats fish per week	1,003	1.97	1.11	0.58	5

Notes: [a]Measured on a scale from 0 (never) to 10 (always);[b] the numeracy test presented subjects with two grids displaying risks of 5 in 10,000 and 10 in 10,000 and asked which one was larger.

4 Econometric Approach

Two types of data emerge from the belief elicitation described above. We directly obtain fractional responses on both the prior and the posterior beliefs about the risk that a foodborne illness is from eating fish. Based on whether the observed belief revisions comply with our definition of consistent updating, we also obtain binary responses.[10] Below, we outline our econometric approach to the analysis of both fractional and binary response data.

4.1 Fractional Responses

The traditional approach to the analysis of fractional responses from the open unit interval is to transform the data to the real line and then apply standard linear regression models. However, such linearized models exhibit a number of unattractive features (Ferrari and Cribari-Neto 2004; Smithson and Verkuilen 2006). The estimated parameters are interpretable in terms of the mean of the transformed rather than the original response. Moreover, linearized models ignore that fractional response data are typically heteroskedastic, i.e. they display less variation at the limits of the unit interval than in its middle range. Lastly, distributions of elicited proportions are often asymmetric causing biased standard errors in small sample estimations.

The double-index beta regression model introduced by Ferrari and Cribari-Neto (2004) and refined by Smithson and Verkuilen (2006) overcomes these problems, making it the preferred statistical approach to analyze data on the open unit interval (0,1).[11] As its name implies, the model has two index func-

[10] One caveat applies to our classification exercise. At the outset we did not know the precautionary behavior of French consumers. The only quantitative information we provided to the subjects was that, on average, French consumers eat fish three times per week (information from the French Ministry of Agriculture, Food and Forestry, www.franceagrimer.fr). In addition, we informed them that older people are more likely to contract a foodborne illness than younger people; that the majority of French consumers do not eat raw fish; and that storing fresh fish for more than three days significantly increases the risk of contracting a foodborne illness. The classification uses only the information about consumer i's deviation from the average frequency of fish consumption and hence that is the only information included in $\Delta\mathbf{D}_i$.

[11] 16 out of our 1,003 subjects stated posterior risk belief of zero. Following Smithson and

tions: one for the mean and one for the precision. To characterize the mean and precision functions as combinations of predictors, Ferrari and Cribari-Neto (2004) re-parameterized the beta distribution to:

$$f(y|\mu,\phi) = \frac{\Gamma(\phi)}{\Gamma(\mu\phi)\Gamma((1-\mu)\phi)} y^{\mu\phi-1}(1-y)^{(1-\mu)\phi-1}, \quad 0 < y < 1,\ 0 < \mu < 1,\ \phi > 0, \tag{4}$$

where $\Gamma(\cdot)$ denotes the gamma function.[12] The beta-distributed response variable y has mean $\mathsf{E}[y] = \mu$ and variance $\mathsf{Var}[y] = \mu(1-\mu)/(1+\phi)$, where ϕ is a precision parameter because, for μ fixed, the larger ϕ the smaller the variance.

Assume now that the fractional responses of our $i = 1, 2, ..., n$ subjects are beta-distributed, i.e. $q_i \sim \mathcal{B}(\mu_i, \phi_i)$. The parameters μ_i and ϕ_i can be mapped onto the real line by appropriate link functions. We use the logit link for the mean function and the log link for the precision function, respectively:

$$g(\mu_i) = \log(\mu_i/(1-\mu_i)) = \beta'\mathbf{X}_i \leftrightarrow \mu_i = \frac{\exp(\beta'\mathbf{X}_i)}{1+\exp(\beta'\mathbf{X}_i)}, \text{ and} \tag{5}$$

$$h(\phi_i) = \log(\phi_i) = \zeta'\mathbf{Z}_i \leftrightarrow \phi_i = \exp(\zeta'\mathbf{Z}_i). \tag{6}$$

Here, $\mathbf{X}_i$ and $\mathbf{Z}_i$ are vectors of covariates associated with the mean and precision parameter of the beta distribution, respectively. The corresponding coefficient vectors β and ζ are estimated by maximum likelihood techniques and statistical inference about their estimates is based on the central limit theorem.

So far, we have assumed that heterogeneity in belief updating can be sufficiently explained by observable characteristics of the subjects. However, latent characteristics may also affect the formation and revision of risk beliefs (Smith et al. 2001). Finite mixture models have become a popular means to capture such unobserved heterogeneity among decision makers (e.g., Andersen

Verkuilen (2006) we used the ad-hoc transformation $q = (\tilde{q}(n-1)+0.5)/n$ (where $\tilde{q}$ is the untransformed fractional response and $n = 1,003$) to restrict all observations to the open unit interval.

[12] The beta distribution is commonly expressed as $f(y|a,b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} y^{a-1}(1-y)^{b-1}, \quad 0 < y < 1$. Equation (4) is obtained by replacing $a = \mu\phi$ and $b = (1-\mu)\phi$.

et al. 2008; Bruhin et al. 2010). The goal of mixture modeling is to estimate the set of parameters in $g(\cdot)$ and $h(\cdot)$ for each of $k = 1, 2, \ldots, K$ latent classes along with their membership function, so that the model predicts to which class each observation most likely belongs (McLachan and Peel 2000).

Although well established in other fields, empirical applications of beta regression mixtures in the social sciences are still rare. Notable exceptions are the studies by Smithson and colleagues (2009; 2011), which present estimates of the beta regression mixture model:

$$m(q|\mathbf{W}, \mathbf{X}, \mathbf{Z}, \psi) = \sum_{k=1}^{K} \pi_k(\mathbf{W}, \alpha_k) f_k(q|\mu(\mathbf{X}, \beta_k), \phi(\mathbf{Z}, \zeta_k)), \tag{7}$$

where the vector $\psi = (\alpha_1, ..., \alpha_K, \beta_1, ..., \beta_K, \zeta_1, ..., \zeta_K)$ collects the class-specific parameters of the discrete mixture density $m(\cdot)$; $f_k(\cdot)$ is the density of the parameterized beta distribution that belongs to the kth latent class; and $\pi_k(\cdot)$ is the corresponding membership function, which depends on a class-specific constant and, possibly, on concomitant variables summarized in the vector $\mathbf{W}$ (Dayton and Macready 1988). The membership function itself is conveniently modeled by the multinomial logit:

$$\pi_k(\mathbf{W}, \alpha_k) = \frac{\exp(\alpha_k' \mathbf{W})}{\sum_{u=1}^{K} \exp(\alpha_u' \mathbf{W})}, \tag{8}$$

whose identification requires normalizing the coefficient vector of class c to zero (i.e., $\alpha_c = 0$).

The mixture model (7) has a non-trivial likelihood function, the maximization of which is cumbersome. We employ the expectation maximization (EM) algorithm (Dempster et al. 1977) as implemented in the R package 'betareg' (version 3.0-4). The EM algorithm is an iterative procedure for maximizing likelihood functions in a missing data setting. In the context of latent class mixture models, it is the class membership of the subjects that is unknown. The EM algorithm iterates between the E-step: evaluation of the expected complete-data log-likelihood given the observed data by fitting equation (8) to each observation; and the M-Step: maximization of the complete-data log-likelihood pertaining to equation (7), using previously derived individual posterior class probabilities as weights.[13] The iteration continues until the EM algorithm converges to a stationary point of the likelihood function (McLachan and Peel 2000). In the empirical application, we set the convergence criterion to 1E-10 and repeated the estimation with 100 random starting values to reduce the risk of running into a local maximum.

4.2 Binary Responses

The most popular models to analyze binary response data are the standard probit and logit, which rest on the assumption that the error distribution of the underlying sorting process has unit variance. This assumption is invalid whenever one group of subjects display greater variability in their response to information than another. On one hand, such heterogeneity in variance across subjects can arise due to uncertainty, misperception, or inattention. On the other hand, subjects belonging to the same observable group might share latent traits that systematically affect the processing of risk-related information.

Consider the non-shaded areas of Figure 1. Belief revisions in these areas are inconsistent unless the subject differs from the average consumer in ways prescribed by our consistency definition. Yet the violations of the Bayesian learning model are qualitatively distinct. Subjects with prior risk beliefs located in areas E and G (D and F) display baseline pessimism (optimism) in the sense of Spinnewijn (2013): without information these subjects perceive their risk to be larger (smaller) than the statistical risk, and upon receiving information they further increase (reduce) their risk estimate. Posterior risk beliefs located in areas D and E (F and G) suggest control pessimism (opti-

[13] In each iteration j, the E-Step calculates a Bayesian update of every subject's posterior probability of belonging to class k:

$$\tau_{ik}(q_i|\psi_i^{(j)}) = \frac{\pi_k(\mathbf{W}, \alpha_k^{(j)}) f_k(q|\mu(\mathbf{X}, \beta_k^{(j)}), \phi(\mathbf{Z}, \zeta_k^{(j)}))}{\sum_{u=1}^{K} \pi_u(\mathbf{W}, \alpha_u^{(j)}) f_u(q|\mu(\mathbf{X}, \beta_u^{(j)}), \phi(\mathbf{Z}, \zeta_u^{(j)}))}.$$

As Bruhin et al. (2010) note, the final posterior probabilities of belonging to class k is a valuable result of the estimation because it provides information on the sharpness of the classification.

mism): subjects believe their personal behavior increases (reduces) the risk above (below) the statistical risk.

It seems unlikely that different inconsistencies in updating share a common error distribution as some subjects will be less certain than others leading to a wider variation in their revised beliefs. A straightforward way to deal with heteroskedasticity in binary response data is to rescale the variance of subgroups based on some predictors of interest (Davidson and MacKinnon 1984). The resulting heteroskedastic binary choice model has the likelihood function:

$$\mathcal{L}\left(\beta, \varphi | \mathbf{X}_i, \mathbf{V}_i\right) = \prod_{i=1}^{n} \left[F\left(\frac{\beta' \mathbf{X}_i}{\exp\left(\varphi' \mathbf{V}_i\right)}\right)\right]^{I_i} \left[1 - F\left(\frac{\beta' \mathbf{X}_i}{\exp\left(\varphi' \mathbf{V}_i\right)}\right)\right]^{1-I_i}, \quad (9)$$

where $F(\cdot)$ denotes the respective cumulative distribution function; I_i indicates inconsistent updating according to the consistency definition; $\mathbf{X}_i$ and $\mathbf{V}_i$ are vectors of covariates associated with the mean and scale, respectively;[14] β and φ are the corresponding coefficient vectors. In the empirical application, we estimate heteroskedastic probit models and include indicators for attitudes toward baseline risk and risk controllability in $\mathbf{V}_i$.

5 Results

We arrange the presentation of our results around the following three questions: Do people update their beliefs in a Bayesian manner? What factors account for heterogeneity in belief updating? And what factors may predict inconsistent belief revisions?

5.1 Observed Belief Updating

The histogram in panel A of Figure 2 illustrates that, before receiving information about the average consumer's risk, subjects were relatively uncertain about their personal risk (mean = 32%, median = 25%). Moreover, we find a

[14] For identification reasons, $\mathbf{V}_i$ cannot contain a constant (Davidson and MacKinnon 1984).

significant spike at 50% suggesting that some subjects had “no idea as to the answer” (Fischhoff and Bruine De Bruin 1999). The provision of risk-related information reduced the perceived risk significantly (mean = 23%, median = 16%) and smoothed out the spike (Figure 2, panel B). However, the average posterior subjective risk belief was still 8 percent points higher than the statistical risk s, suggesting that on the sample level subjects displayed some control pessimism (Spinnewijn 2013).

Figure 2: Histograms of observed belief revisions

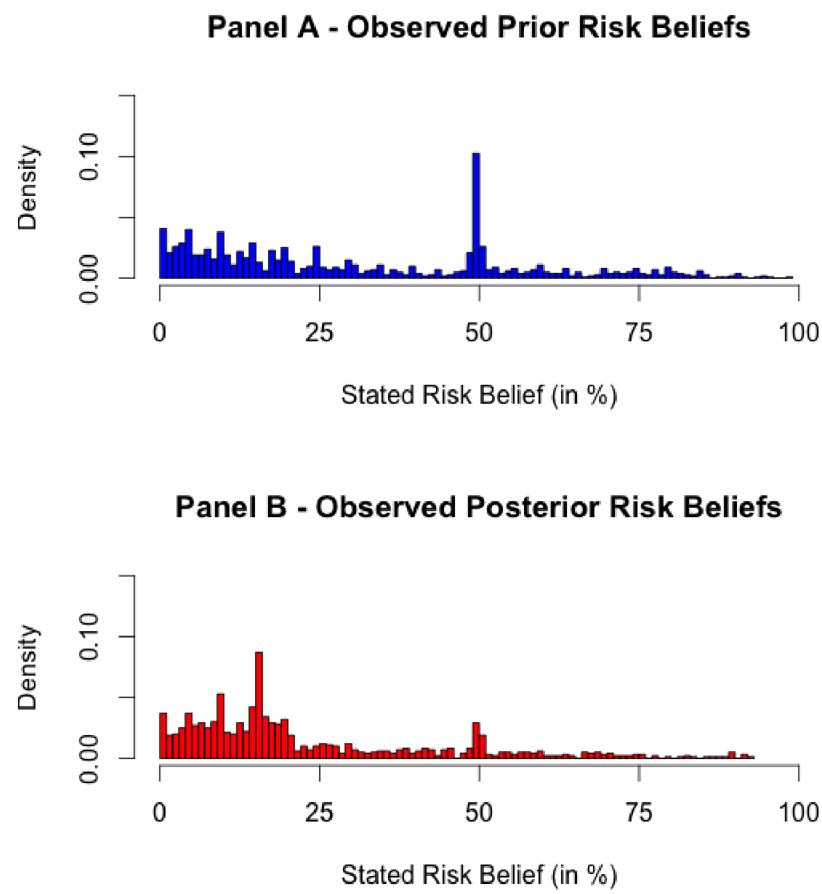

Figure 3: Landscape of observed belief revisions

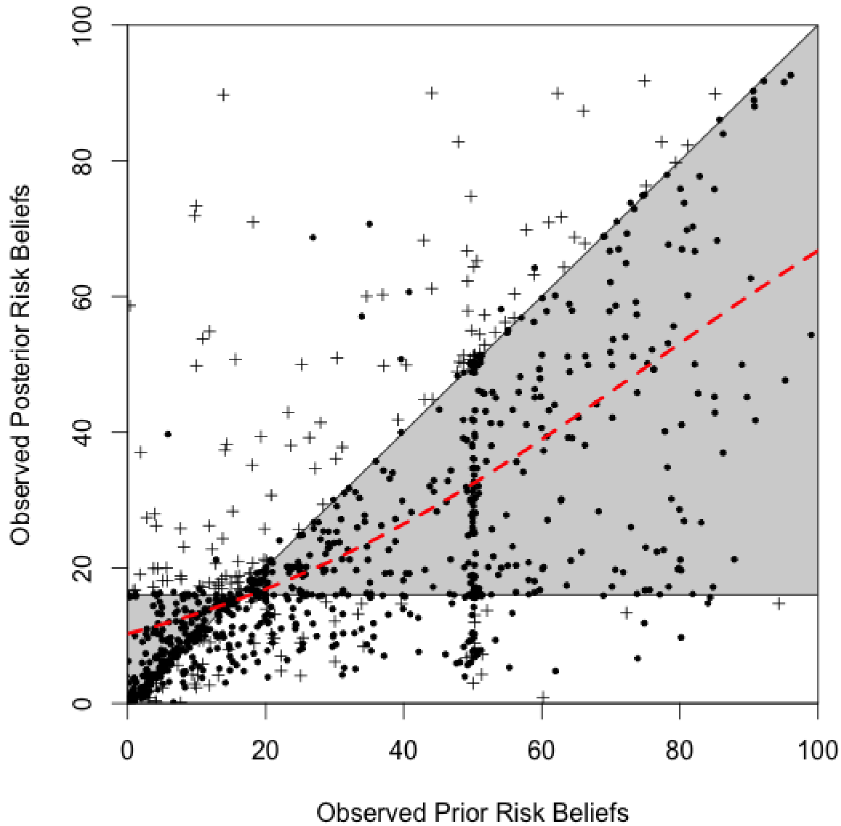

Notes: Points (cross marks) denote observations that do (not) comply with the consistency definition. The red-dotted line gives the fit of the most basic beta regression model.

Figure 3 plots the posterior against the prior risk beliefs. The vast majority of subjects (81%) updated their risk belief in response to the information provided. The dots on the horizontal line in Figure 3 represent 85 subjects whose revised beliefs perfectly matched the statistical risk (i.e., $\xi^* = 1$). Similarly, the dots on the diagonal line represent 186 subjects who dismissed the information altogether (i.e., $\xi^* = 0$). The crosses mark 175 subjects who violate the conditions of the consistency definition. The dashed line represents the best fit obtained from a simple beta regression of q_i on p_i.[15] It confirms

[15] Since s is constant across subjects, its effect cannot be separated from the regression intercept. Normalizing the relative precision weights in equation (1) to sum to one, the estimated precision weight attached to the information is $\hat{\xi}^* = 1 - \hat{\gamma}^* = 0.493$.

that on average subjects adjusted their beliefs upward (downward) if their prior was smaller (larger) than the statistical risk.

5.2 Observed Heterogeneity in Belief Updating

Next we turn our attention to heterogeneity in the revision of risk beliefs. It is widely accepted that personal characteristics and world views mediate the perception of risks (Slovic 1999). The same factors may also affect how new information is processed to update beliefs about these risks. We explore heterogeneity in the updating process using the workhorse model (3). Because risk beliefs are bounded to the unit interval, we employ the double-index beta regression model for this analysis. In Table 2 we report on two distinct specifications of the beta regression model. The naïve model ignores that factors which affect the updating of risk beliefs might also affect the formation of prior beliefs and/or the precautionary effort expended by the consumer. With respect to equation (3), the naïve model thus hypothesizes: $\beta_4 = \beta_5 = \beta_6 = \beta_7 = 0$. In contrast, the extended model includes all two- and three-way interactions between the socioeconomic characteristics of the subject, the precautionary effort, and the prior risk belief. In both specifications, we model the precision in belief updating as a function of age, gender, educational attainment, and whether or not the subject passed a simple numeracy test. We hypothesize that these factors are related to the numerical skills needed to process risk-related information (Peters et al. 2006).

We discuss the major findings of the beta regression analysis in terms of first differences in predicted posterior risk beliefs (reported in Table 2) because the interpretation of interaction effects in generalized linear models is complicated by the link function (Ai and Norton 2003). Figure 4 plots the predicted posterior beliefs for both specifications against the observed priors holding all other variables fixed at sample means. The plot confirms the finding that, on average, subjects substantially reduced their prior risk beliefs upon receiving the risk-related information. The dashed plots show counterfactual

Table 2: Observed updating: results of beta regression models

	Naïve Model						Extended Model					
	Est.		Std. Error	Data Range	1^{st}Diff.	95% Conf. Int.[a]	Est.		Std. Error	Data Range	1^{st}Diff.	95% Conf. Int.[b]
(Intercept)	-1.566	***	0.263				-1.524	***	0.441			
PRIOR	2.720	***	0.135	[0.0; 1.0]	0.526	[0.480; 0.572]	2.305	*	1.178	[0.0; 1.0]	0.511	[0.455; 0.547]
HEALTH STATUS	-0.012		0.017	[0.0; 10.0]	-0.020	[-0.076; 0.032]	0.028		0.023	[0.0; 10.0]	-0.041	[-0.112; 0.021]
PREGNANCY	0.127		0.083	[0.0; 1.0]	0.021	[-0.007; 0.048]	-0.013		0.150	[0.0; 1.0]	0.023	[-0.011; 0.060]
MALE	-0.199	***	0.056	[0.0; 1.0]	-0.033	[-0.051; -0.014]	0.014		0.103	[0.0; 1.0]	-0.028	[-0.052;- 0.007]
KIDS	0.031		0.062	[0.0; 1.0]	0.005	[-0.016; 0.025]	-0.225	**	0.107	[0.0; 1.0]	-0.001	[-0.024;- 0.023]
AGE	0.004		0.003	[18.0; 80.0]	0.041	[-0.015; 0.095]	-0.009	**	0.004	[18.0; 80.0]	0.040	[-0.024; 0.099]
INCOME[b]	-0.022		0.027	[0.25; 5.0]	-0.017	[-0.057; 0.024]	0.001		0.047	[0.25; 5.0]	-0.012	[-0.062; 0.031]
EDUCATION	-0.028	***	0.009	[4.0; 17.0]	-0.058	[-0.095; -0.022]	-0.043	***	0.016	[4.0; 17.0]	-0.072	[-0.121; -0.029]
NUMERACY	-0.209	**	0.104	[0.0; 1.0]	-0.036	[-0.072;- 0.001]	0.054		0.316	[0.0; 1.0]	0.008	[-0.059; 0.052]
SUBJECTIVE RISK	0.136	**	0.058	[0.0; 1.0]	0.022	[0.004; 0.042]	0.051		0.097	[0.0; 1.0]	0.022	[0.001; 0.043]
SAFETY CONCERNS	0.039		0.052	[0.0; 1.0]	0.006	[-0.011; 0.024]	0.193	**	0.084	[0.0; 1.0]	0.003	[-0.016; 0.022]
MEALS[c]	0.014		0.025	[0.57; 5.0]	0.010	[-0.026; 0.047]	1.171	***	0.326	[0.57; 5.0]	0.029	[-0.015; 0.071]
RAW FISH[c]	0.028		0.018	[0.0; 10.0]	0.048	[-0.010; 0.111]	-0.284		0.276	[0.0; 10.0]	0.073	[0.015; 0.160]
WASH HANDS[c]	-0.016	**	0.007	[0.0; 10.0]	-0.027	[-0.051; -0.003]	-0.037		0.107	[0.0; 10.0]	-0.034	[-0.060;-0.005]
STORE FISH[c]	0.028	**	0.011	[0.0; 10.0]	0.047	[0.008; 0.086]	0.122		0.180	[0.0; 10.0]	0.039	[-0.001; 0.110]
PREPARE FISH[c]	-0.014	**	0.007	[0.0; 10.0]	-0.024	[-0.046;-0.001]	-0.046		0.111	[0.0; 10.0]	-0.029	[-0.058; -0.004]
(ϕ)_(Intercept)	1.654	***	0.345				1.040	***	0.316			
(ϕ)_MALE	0.315	***	0.113				0.363	***	0.104			
(ϕ)_AGE	-0.010	**	0.004				-0.004		0.004			
(ϕ)_EDUCATION	0.024		0.016				0.055	***	0.014			
(ϕ)_NUMERACY	0.319	**	0.154				0.299	**	0.147			
Interaction effects	excluded						included ($\chi^2 = 148.7$,115 Df)					
Observations	1,003						1,003					
Log-likelihood	-809.7						-884.1					
AIC	-1,575.4						-1,494.2					
BIC	-1,467.4						-821.4					

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a]based on 1,000 bootstrap replications; [b] missing values approximated by the observed sample mean; [c] mean-centered variables.

predictions for particularly risky and precautionary behaviors.[16] While the model predictions are almost identical for the most precautionary consumer, they differ somewhat for the most risk-taking consumer. For prior risk beliefs smaller (larger) than $p = 35\%$, predictions based on the naïve model are distinctly more (less) control pessimistic than those based on the extended model.

Figure 4: Landscape of predicted belief revisions

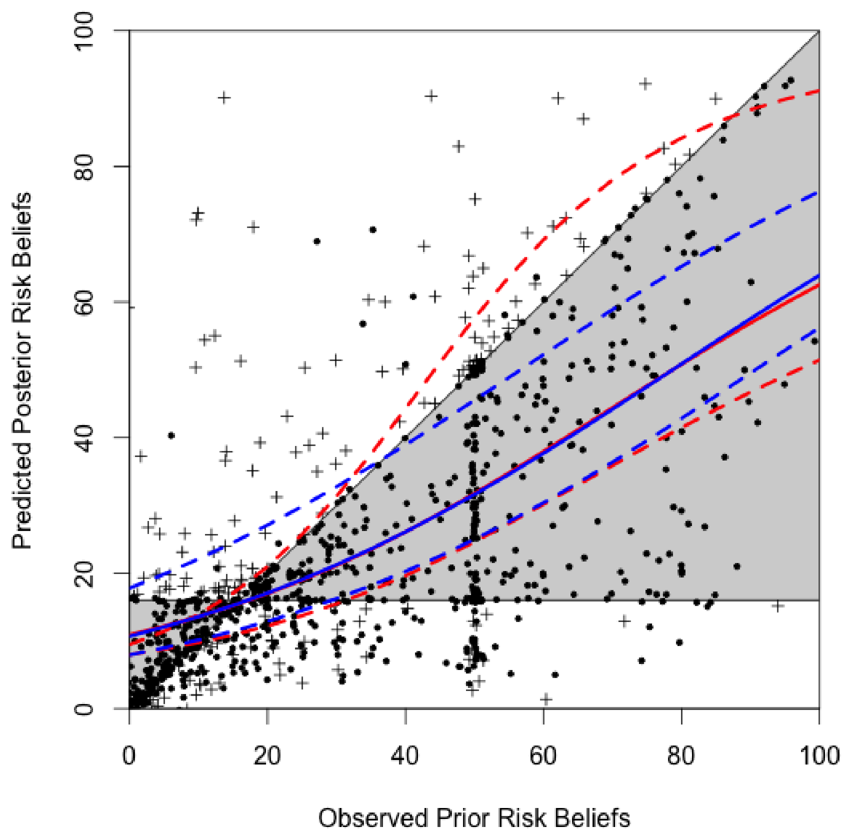

Notes: Points (cross marks) denote observations that do (not) comply with the consistency definition. The blue and red fits are based on the naïve and the extended model (Table 2), respectively; the dotted fits represent hypothetical consumers who are most and least precautious, respectively (see the main text for details).

[16]For the counterfactuals, we replaced the reported values of the variables `MEALS`, `RAW FISH`, `WASH HANDS`, `STORE FISH`, and `PREPARE FISH` by their observed maximum and minimum values, respectively.

Figure 4 suggests that precautionary effort plays an important role in understanding the updating of subjective risk beliefs. This conclusion is supported by the coefficient estimates of the naïve model (column 1 in Table 2). In particular, we find that subjects who take more (less) precaution than the average consumer stated significantly lower (higher) posterior risk beliefs. First differences indicate that heterogeneity in each of the recorded precautionary behaviors may account for differences in posterior risk beliefs of 2.4 to 4.8 percent points.

Perceived vulnerability matters less than expected. The only significant predictor is the `SUBJECTIVE RISK` of suffering a foodborne illness, which was associated with a 2.2 percentage point larger posterior risk belief among subjects who reported they bear a high risk. In agreement with the gender effect observed in many risk perception studies (Slovic 1999), `MALE` subjects had a 3.3 percentage point smaller posterior risk belief than female subjects. Posterior risk beliefs were negatively associated with `EDUCATION`, decreasing by about 0.4 percent points per additional year of schooling. Notably, other indicators related to risk literacy such as the `NUMERACY` test and the `AGE` of the subject had no statistically significant association with the revision of risk beliefs.

Although the results of the naïve model are consistent with what one would intuitively expect, the results of the extended model (column 2 in Table 2) call for a more cautious interpretation. Indeed, once we account for possible interaction effects between prior risk beliefs and the characteristics of the subject, the main effects related to the precautionary behaviors are no longer significant. Yet the first differences related to precaution are of comparable size to those of the naïve model, suggesting that precautionary behaviors affect the formation of the prior risk belief rather than the revision process. We conclude that ignoring these interdependencies—as we do with the naïve model—leads to overestimating the impact of precaution on the belief revision process.

By controlling for potential confounding effects, the extended model reveals an important main effect that is masked in the naïve model. Subjects increased their posterior risk belief by 2.9 percent points per weekly fish `MEAL`, meaning that there is significant response to the quantitative part of the information provided. Even after controlling for the impact of education on the formation of the prior risk belief, more-educated subjects state lower posterior beliefs (0.6 percent points per year). As the interaction effect between the `PRIOR` and the years of `EDUCATION` is statistically not different from zero (not reported in Table 2), the latter finding might be interpreted as implying that more-educated subjects put more weight on the information if their prior risk belief were higher than the statistical risk. Again, we do not observe that indicators of numerical skills other than education are associated with the updating of risk beliefs.

In both specifications the modeled precision in belief updating decreases with `AGE` and is higher among `MALE` subjects. Precision also increases with the years of attained `EDUCATION` and is significantly higher for subjects who passed the `NUMERACY` test, indicating that young, more-educated males responded far more homogeneously to the information provided than older, less-educated females.

5.3 Unobserved Heterogeneity in Belief Updating

So far we have implicitly assumed that heterogeneity in belief updating can be captured by observable characteristics of the subjects. Indeed, the analysis above identified a number of predictors that explain the updating of risk beliefs. There is, however, good reason to believe that unobserved characteristics also affect the revision process (Hogarth and Einhorn 1992). The intuition is that response to information varies from one group of subjects to another, but the corresponding group memberships are unobserved.

Table 3 presents our preferred beta regression mixture model, which presumes there exist four latent classes of belief updaters.[17] The belief updating

[17] We explored various beta regression mixtures with up to six latent classes.Four classes yielded the best results in terms of common information criteria (AIC, BIC, ICL-BIC) throughout the analysis. The appendix provides a summary of these goodness-of-fit measures and reports specification and robustness tests. While the membership composition varies somewhat across the different model specifications, the thrust of the results remains unchanged.

Table 3: Unobserved updating: results of beta regression mixture modeling

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.032	***	0.128	-2.781	***	0.085	-2.404	***	0.024	-2.126	***	0.106
`PRIOR`	3.032	***	0.319	4.464	***	0.181	4.808	***	0.005	1.041	***	0.237
`MEALS`[a]	0.034		0.107	0.013		0.054	0.038	**	0.020	-0.206	**	0.082
`PRIOR` × `MEALS`[a]	-0.459	*	0.239	0.162		0.106	-0.101	**	0.041	0.672	***	0.156
(ϕ)_(Intercept)	1.216	***	0.130	3.546	***	0.153	7.272	***	0.326	3.577	***	0.195
(π)_(Intercept)	fixed to 0			-1.488		1.158	-0.661		0.995	-2.872	*	1.610
(π)_`MALE`	fixed to 0			0.465		0.341	0.055		0.350	0.350		0.367
(π)_`AGE`	fixed to 0			-0.021		0.013	-0.022	*	0.013	-0.029	**	0.014
(π)_`EDUCATION`	fixed to 0			0.099	**	0.047	0.065		0.049	0.189	***	0.061
(π)_`NUMERACY`	fixed to 0			1.467	**	0.624	0.238		0.487	1.743		1.117
Membership size[b]	131			420			176			276		
Observations	1,003											
Log-likelihood	-1,003.0											
AIC	-1,936.0											
BIC	-1,764.2											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k.

function μ_k reflects that `MEAL` is the only predictor related to precautionary effort that seems to have direct impact on the belief revision process (see the extended model in Table 2). The membership function π_k includes gender, age, educational attainment, and numeracy as concomitant variables.[18] As the identification of the latent classes is co-determined by these indicators, the model includes a unique precision term ϕ_k for each latent class k.

The mixture model clusters the observed relationship between prior and posterior beliefs into four latent classes. This classification is based on the highest class membership probability at convergence, i.e. subject i belongs to latent class k iff $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \, \forall u$. Each of the classes represents a distinct pattern of belief revision. Figure 5 shows the class-specific updating functions and the corresponding class membership probabilities at convergence (whereby brighter colors indicate a higher probability of belonging to class k).

The purple dots represent 176 subjects who ignored the information provided and refused to substantially revise their risk beliefs. We hence refer to this class as *information refuseniks*. The blue dots represent the class of *belief adopters*, comprising 276 subjects whose revised risk beliefs are close to the statistical risk s. The green dots represent the class of *Bayesian learners*, which includes 420 subjects who combined their prior risk belief and the information provided in a manner consistent with the Bayesian learning hypothesis to form a new posterior risk belief. Lastly, the red dots represent 131 subjects whose updating behavior largely overlaps with the belief revisions of subjects whom we classified as *inconsistent updaters* according to the consistency definition.[19]

The results displayed Figure 5 warrant some remarks. First, the classification obtained from the mixture model is far from crisp; i.e., for most subjects the maximum probability of membership in any class is much less than one. We find that $\tau_{ik}^{(*)} \geq \frac{1}{2}$ holds for 80% of the observations, $\tau_{ik}^{(*)} \geq \frac{2}{3}$ holds for 45% of the observations, and $\tau_{ik}^{(*)} \geq \frac{9}{10}$ holds for 14% of the observations. Such a

[18] We keep predictors of the updating function and the membership function deliberately separated as otherwise the identifying assumption of local independence might be violated (McLachlan and Peel 2000).

[19] There is a 84% overlap of the classifications obtained by mixture modeling and based on the consistency definition.

noisy classification is not unusual in empirical applications of mixture models and reflects the inverse relationship between the number of latent classes and the discriminatory power of the mixture model.

Figure 5: Latent patterns of belief updating

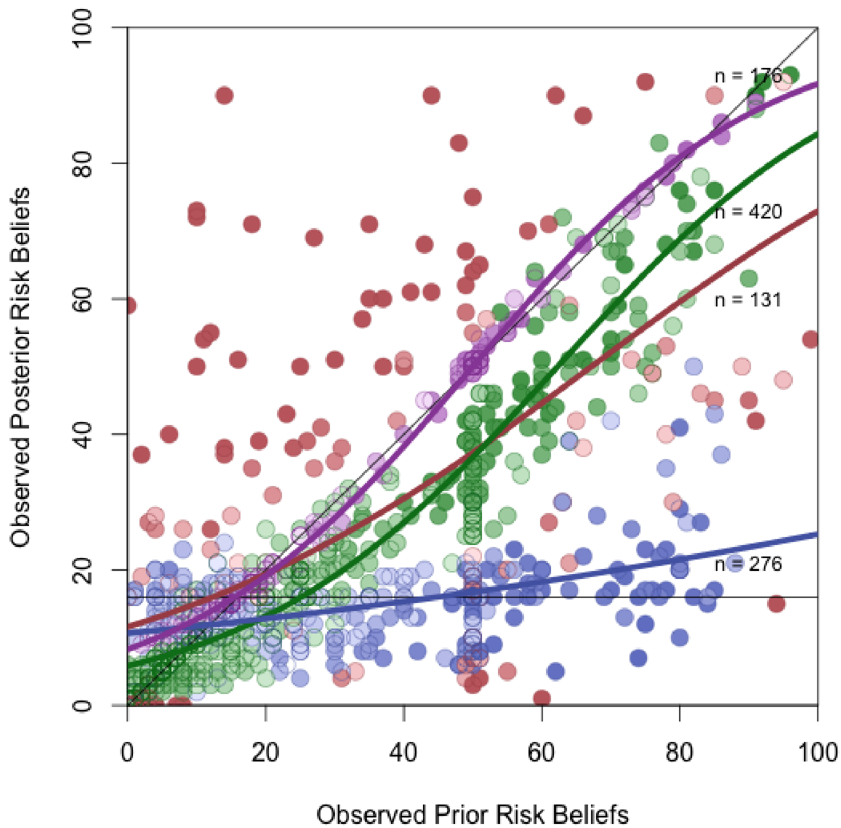

Notes: Colors reflect class memberships; brighter colors indicate higher class membership probability at convergence; lines represent class-specific beta regression fits.

Second, we find that the precision in updating among inconsistent updaters is up to six times smaller than among members of the other classes. This highlights the element of randomness in the belief revision of inconsistent updaters. Bayesian learners and belief adopters have almost the same estimated precision, suggesting that there is no sharp boundary between the two updating strategies. In contrast, information refuseniks have a very high precision estimate. The posterior beliefs of this class are close to their prior beliefs, whether low or high.

Third, the updating functions of the four classes are distinctly different. Inconsistent updaters who eat fish more (less) often than the average consumer in the sample reduce (increase) their posterior risk belief in response to the information. This highlights the similarity between the class of inconsistent updaters as classified by the mixture model and violators of the consistency definition. Members of other latent classes increase (reduce) their posterior risk beliefs if their fish consumption is higher (lower) than average, with information refuseniks giving least weight and belief adopters giving most weight to their consumption.

Fourth, estimates of the class membership function indicate that numeracy, age, and educational attainment influence the likelihood of belonging to a specific class k. In particular, the probability to be classified as inconsistent updater increases with `AGE`, is higher among subjects who failed the `NUMERACY` test and had attained less `EDUCATION`. Apart from their age, the class of information refuseniks is not much different from the class of inconsistent updaters, suggesting that the refusal to update might be a strategy of coping with information one cannot process.

5.4 Inconsistent Belief Updating

Concomitant-variable latent class models are a powerful method for classifying observations based on both unobservable and observable characteristics (Dayton and Macready 1988). One drawback, however, is that we estimate the impact of the concomitant variables on the membership function of each of the $K-1$ classes (using class c as the baseline). We obtain a valid between-class comparison, but as soon as there are more than two classes it is difficult to compare the effect of the concomitant variables on one particular group of interest to the common effect on *all* other classes.

Let us put the interpretation issue into the context of our analysis. If we seek to better understand inconsistent updating, it seems natural to compare inconsistent updaters to consistent updaters rather than to different types of consistent updaters. Moreover, such a comparison allows us to explore similarities and differences between the parametric classification based on the mixture modeling and the non-parametric classification of the consistency definition. Table 4 reports the estimates of two heteroskedastic probit models that link inconsistent updating as established by the parametric classification (PC) of the beta regression mixture model and by the non-parametric classification (NPC) of the consistent belief updating definition to indicators of consumption, precautionary effort, and socio-economic factors. In both specifications, we test for heteroskedasticity across subjects with different attitudes toward the baseline risk and risk controllability.

We discuss significant predictors of inconsistent updating in terms of their estimated marginal effects at sample means.[20] These can be readily interpreted as the expected change in the probability of being an inconsistent updater conditional on a change in the value of the particular predictor. Under both classifications, the probability of being classified as inconsistent updater increases by 6.3 to 8.9 percentage points if the subject lives with young `KIDS` and it decreases by 0.2 to 0.3 percentage points for each year of `AGE`. Although only significant in the NPC model, we find indication of a gender effect with `MALE` subjects being 2.7 to 5 percentage points less likely to inconsistently revise their risk beliefs. In both specifications eating fish more often than the average consumer increases the likelihood of being classified as inconsistent. However, the consumption effect is much stronger in the NPC model (9.4 percentage points for each additional `MEAL`) than in the PC model (1.7 percentage points for each additional `MEAL`), where in addition `RAWFISH` consumption raises the likelihood of being classified as inconsistent (by 1.1 percentage point per additional `RAWFISH` meal).

Attitudes toward baseline risk and controllability have a large and significant effect on the likelihood to inconsistently revising risk beliefs. Among subjects exhibiting `BASELINE PESSIMISM` the likelihood of inconsistent up-

[20]For the heteroskedastic probit model, the marginal effect of a change in variable w is $\frac{\partial \Pr(I=1|\mathbf{X},\mathbf{V})}{\partial w} = \phi\left[\frac{\beta'\mathbf{X}}{\exp(\varphi'\mathbf{V})}\right]\frac{\beta_w-(\beta'\mathbf{X})\varphi_w}{\exp(\varphi'\mathbf{V})}$, with only the first (second) term on the RHS applying if w appears only in $\mathbf{X}$ ($\mathbf{V}$).

Table 4: Inconsistent updating: results of heteroskedastic probit models

	Parametric Classification Model					Non-Parametric Classification Model				
	Est.		Std. Error	Marginal Effects	95% Conf. Interval	Est.		Std. Error	Marginal Effects	95% Conf. Interval
(Intercept)	0.403		0.487			-2.284	***	0.768		
HEALTH STATUS	-0.032		0.037	-0.005	[-0.016; 0.006]	-0.001		0.044	0.000	[-0.013; 0.013]
PREGNANCY	-0.074		0.247	-0.011	[-0.081; 0.059]	-0.195		0.349	-0.028	[-0.122; 0.065]
MALE	-0.179		0.138	-0.028	[-0.069;- 0.014]	-0.329	**	0.163	-0.050	[-0.099;- 0.002]
KIDS	0.368	**	0.161	0.063	[0.019; 0.121]	0.538	**	0.208	0.089	[0.019; 0.158]
AGE	0.023	***	0.006	0.004	[0.002; 0.005]	0.015	*	0.008	0.002	[0.000; 0.005]
INCOME[a]	-0.054		0.061	-0.008	[-0.027; 0.010]	0.083		0.080	0.013	[-0.011; 0.036]
EDUCATION	-0.092	***	0.018	-0.014	[-0.020; -0.009]	-0.017		0.025	-0.003	[-0.010; 0.005]
NUMERACY	-1.183	***	0.211	-0.303	[-0.430; -0.176]	-0.024		0.366	-0.004	[-0.115; 0.108]
SUBJECTIVE RISK	-0.124		0.141	-0.019	[-0.062; 0.023]	0.022		0.173	0.003	[-0.049; 0.055]
SAFETY CONCERNS	0.048		0.126	0.007	[-0.031; 0.046]	0.194		0.156	0.030	[-0.018; 0.077]
MEALS[b]	0.108	*	0.059	0.017	[-0.002; 0.035]	0.620	***	0.088	0.095	[0.071; 0.119]
RAW FISH[b]	0.074	**	0.035	0.011	[0.001; 0.022]	-0.050		0.054	-0.008	[-0.024; 0.008]
WASH HANDS[b]	0.012		0.017	-0.004	[-0.013;0.005]	-0.002		0.020	-0.012	[-0.027;0.003]
STORE FISH[b]	-0.026		0.031	0.002	[-0.003; 0.007]	-0.080		0.052	0.000	[-0.006; 0.006]
PREPARE FISH[b]	0.015		0.017	0.002	[-0.003; 0.007]	-0.011		0.020	-0.002	[-0.008; 0.004]
(σ)_BASELINE PESSIMISM	-0.266	**	0.119	-0.059	[-0.111; -0.008]	-0.341	***	0.125	-0.082	[-0.139; -0.024]
(σ)_CONTROL PESSIMISM	0.490	***	0.120	0.106	[0.058; 0.154]	1.355	***	0.174	0.282	[0.227; 0.337]
Observations	1,003					1,003				
Log-likelihood	-316.7					-398.6				
AIC	-597.5					-761.3				
BIC	-509.1					-672.9				

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$;[a] missing values approximated by the observed sample mean; [b] mean-centered variable.

dating was 5.9 to 8.2 percentage points lower than among baseline-optimistic subjects, while CONTROL PESSIMISM increased the likelihood by 10.6 percentage points (PC model) and 28.2 percentage points (NC model), respectively. Coherent with the results of the beta regression mixture model, we find that in the PC model EDUCATION and NUMERACY have a large effect on the consistency of belief revisions. Indeed, the likelihood of being classified as inconsistent falls by 1.4 percentage points with each additional year of schooling and it is 30.2 percentage points lower among subjects who passed the NUMERACY test. The NC model does not reflect such strong cognitive effects, hinting at the fact that differences in the frequency of fish consumption alone might be insufficient to classify the updating behavior.[21]

6 Discussion

Providing relevant and accurate information to the public is a central aspect of information-based health policies, but it does not suffice for the welfare assessment of such policies. Policy makers also need to know how (if at all) people respond to the provided information. Below, we discuss the key findings of our study both on the individual and the aggregate level of belief revisions.

6.1 Individual Response to Information

The key question related to the efficacy of information-based health policies is who is going to respond by how much to the information provided (Magat and Viscusi 1992). We have examined patterns of belief revision including, in particular, apparent heterogeneity. Surprisingly little of this heterogeneity is explained by the precautionary effort subjects made to reduce the risk of contracting a foodborne illness. Rather, we find that risk perceptions increased with higher consumption of fish and concerns about seafood safety. This finding is consistent with the risk-as-feeling hypothesis (Loewenstein et

[21] This is not a shortcoming of the consistent updating definition per se, but relates to our limited knowledge of $\mathbf{\Delta D}_i$.

al. 2001), which postulates that responses to risk information result in part from feelings such as worry, fear, dread, or anxiety that arise at the time of belief updating. These feelings may exert a reciprocal influence on the cognitive evaluation of risk-related information, suppressing consideration of objectively risk-increasing or risk-reducing factors.

Recent research by Peters et al. (2006) finds that people differ in the degree to which they process risks cognitively versus affectively. In particular, their research indicates that highly numerate people draw more meaning from probabilities, frequencies, and other numerical comparisons than the less numerate do. In consequence, numerical risk information provides less meaning to individuals with lower numerical skills and is therefore less helpful in updating risk beliefs. Moreover, Peters and colleagues found that less numerate individuals are more prone to respond to irrelevant information suggesting that numerical information may even distort belief revisions. These findings on the ability to process numerical information support the results reported in Tables 3 and 4. Less educated, less numerate, and older subjects were much more likely to inconsistently update their risk beliefs.

6.2 Aggregate Response to Information

Even at the aggregate level, some noteworthy observations can be made on the response to the risk information. Unlike a vast number of psychological studies (Harris and Hahn 2011) we do not find unrealistic optimism about future life events. On average, our subjects were pessimistic when they formed their prior risk beliefs ($\mathsf{E}(p) = 32\%$) and remained slightly pessimistic ($\mathsf{E}[q] = 23\%$) upon receiving information about the statistical risk of contracting a fishborne illness and the possible means to control this risk. As the sample was constructed to be representative of French fish consumers, we would expect $\mathsf{E}[q] \approx s$ if subjects were Bayesians. A Mann-Whitney test clearly rejects the hypothesized equality.

As the belief revision protocol was embedded in a larger survey on fish consumption, there might have been a salience effect at play. Recent and unusual events are more memorable and people therefore tend to draw on them when reasoning about experiencing similar events in the future (Gilbert and Wilson 2007). Yet the relative nature of the belief revision task emphasized both the risk of contracting a foodborne illness from fish as well as from other foods, making the salience hypothesis less plausible. Another possible explanation for control pessimism at the aggregate level is alarmist reactions to risk information (Viscusi 1997): people focus on worst-case scenarios when they update risk beliefs. For instance, some subjects might have believed that the eating of raw fish would drastically increase their risk compared to the average consumer. While this would justify higher prior risk beliefs, it is hard to reconcile with higher posterior risk once we control for precautionary behavior. Lastly, we note that the sample median posterior risk belief was equal to the statistical risk ($\mathsf{M}[q] = 16\%$) suggesting that mean control pessimism could be driven by outliers.[22]

To explore the last explanation in more detail we compared consistent to inconsistent updaters, using both the parametric and non-parametric classification of inconsistent updating. We find a stark information effect among consistent updaters, but no effect at all among the inconsistent updaters.[23] This suggests that, as a group, inconsistent updaters could not infer much from the risk-related information. Yet since inconsistent subjects were about twice as often control-pessimistic as they were control-optimistic, we can also reject that they just made random belief updates. But what is it that drives inconsistent belief revisions? The results obtained from both the mixture modeling and the heterogeneous probit analyses confirm the individual-level finding that inconsistent updaters tend to be older, less educated, and less likely to have passed the simple numeracy test than members of the other la-

[22]The sample mean is more affected by high posterior risk beliefs than it is affected by low ones simply because the statistical risk, that is the population mean (which should equal the sample mean for a representative sample), is smaller than 50%.

[23]This observation holds under both the non-parametric classification based on the consistency definition: ($\mathsf{E}[p|I_{\mathrm{NPC}} = 0] = 28.8\%, \mathsf{E}[q|I_{\mathrm{NPC}} = 0] = 21.4\%$) vs ($\mathsf{E}[p|I_{\mathrm{NPC}} = 1] = 32.1\%, \mathsf{E}[q|I_{\mathrm{NPC}} = 1] = 32.3\%$) as well as the parametric classification based on the mixture modeling: ($\mathsf{E}[p|I_{\mathrm{PC}} = 0] = 31.1\%, \mathsf{E}[q|I_{\mathrm{PC}} = 0] = 21.6\%$) vs ($\mathsf{E}[p|I_{\mathrm{PC}} = 1] = 34.6\%, \mathsf{E}[q|I_{\mathrm{PC}} = 1] = 34.4\%$).

tent classes (albeit we do not find statistically significant differences between inconsistent updaters and information refuseniks). This highlights that specific groups within society are particularly prone to misunderstand health risk communications and that addressing these groups is indeed a challenge.

Heterogeneity in belief revision suggests that risk communication methods should be tailored to specific groups of recipients to obtain the optimal effects. Individuals who can process numerical risk information and update their prior beliefs in a consistent fashion may be well-served by providing such statistical information; in contrast, individuals who cannot consistently process such information may be better served by alternative risk-communication messages, including perhaps messages that direct them to take specific actions (e.g., avoid raw fish, keep fish refrigerated until preparation, etc.). Individuals who fail to update their prior beliefs (information refuseniks) may already be well-informed about risks of consuming fish and other foods so that the information provided in the survey has negligible value; alternatively, they may distrust the source of the information and so alternative sources may be more influential.

Although we have identified substantial heterogeneity in patterns of risk belief updating, we have only limited ability to predict how any individual will update. On average, individuals who are more numerate, highly-educated, and younger are better able to process numerical risk information and to update their prior beliefs in a consistent way, but these are weak predictors and conceal much variation. It may be useful in future work to identify more accurate methods for predicting how different individuals will respond to risk information, and what forms of communication are most effective for different individuals.

7 Conclusion

Economists typically employ the Bayesian learning model to predict the success of information-based policies. More often than not, it is presumed that people perfectly update their beliefs. Yet outcomes of information campaigns depend crucially on people's actual behavior, which in turn is contingent upon whether the disseminated information is taken up and processed in the expected way. In this paper we have explored the revision of risk beliefs in the realm of food safety. We find evidence for four patterns or strategies of updating subjective beliefs about the risk of contracting a foodborne illness. Our analysis suggests that the differences between these strategies are not explained by individual precautionary behavior, but are mainly related to indicators of numeracy. Because information campaigns are often targeted at demographic groups with limited numerical skills, policy makers should seek to communicate statistical information in the most accessible way possible—perhaps even at the cost of simplifying—and consider to tailor information to account for the fact that individuals may differ in the type of assistance they need in making health-related choices. Even if policy makers do so, we cannot anticipate everybody to correctly understand and follow public health advisories. Hence, deficiency in risk numeracy has to be taken into account when predicting the outcome of health information policies.

Appendix

Figure 6: Histogram of class membership probabilities at convergence

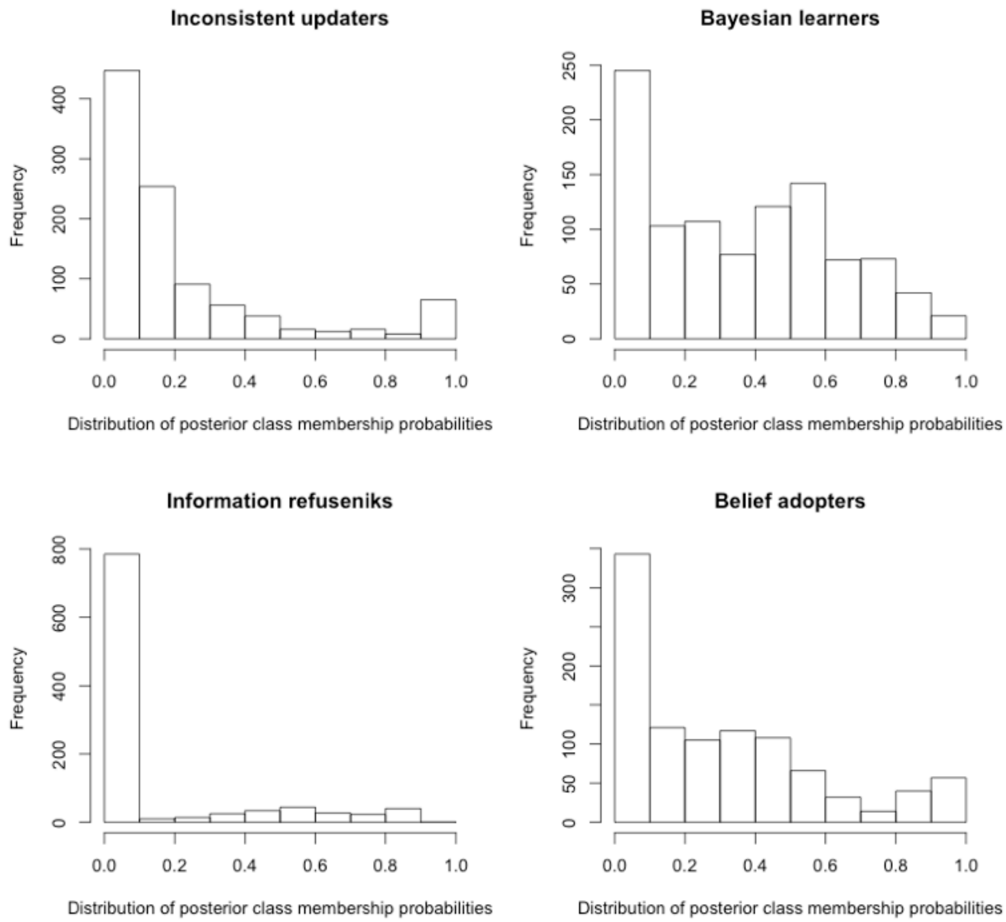

Notes: The histograms correspond to the classes identified by the main beta regression mixture model reported on in Table 3.

Table 5: Specification test: main beta regression mixture model w/o concomitant variables

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.032	***	0.128	-2.781	***	0.085	-2.404	***	0.024	-2.126	***	0.106
PRIOR	3.032	***	0.319	4.464	***	0.181	4.808	***	0.005	1.041	***	0.237
MEALS[a]	0.034		0.107	0.013		0.054	0.038	**	0.020	-0.206	**	0.082
PRIOR × MEALS[a]	-0.459	*	0.239	0.162		0.106	-0.101	**	0.041	0.672	***	0.156
(ϕ)_(Intercept)	1.216	***	0.130	3.546	***	0.153	7.272	***	0.326	3.577	***	0.195
(π)_(Intercept)	fixed to 0			-1.488		1.158	-0.661		0.995	-2.872	*	1.610
Membership size[b]	122			446			182			253		
Observations	1,003											
Log-likelihood	-987.1											
AIC	-1,928.1											
BIC	-1,799.8											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k.

Table 6: Robustness test: main beta regression mixture model, restricted data set

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.032	***	0.128	-2.781	***	0.085	-2.404	***	0.024	-2.126	***	0.106
`PRIOR`	3.032	***	0.319	4.464	***	0.181	4.808	***	0.005	1.041	***	0.237
`MEALS`[a]	0.034		0.107	0.013		0.054	0.038	**	0.020	-0.206	**	0.082
`PRIOR` × `MEALS`[a]	-0.459	*	0.239	0.162		0.106	-0.101	**	0.041	0.672	***	0.156
(ϕ)_(Intercept)	1.216	***	0.130	3.546	***	0.153	7.272	***	0.326	3.577	***	0.195
(π)_(Intercept)	fixed to 0			-1.488		1.158	-0.661		0.995	-2.872	*	1.610
(π)_`MALE`	fixed to 0			0.465		0.341	0.055		0.350	0.350		0.367
(π)_`AGE`	fixed to 0			-0.021		0.013	-0.022	*	0.013	-0.029	**	0.014
(π)_`EDUCATION`	fixed to 0			0.099	**	0.047	0.065		0.049	0.189	***	0.061
(π)_`NUMERACY`	fixed to 0			1.467	**	0.624	0.238		0.487	1.743		1.117
Membership size[b]	131			420			176			276		
Observations	1,003											
Log-likelihood	-1,003.0											
AIC	-1,936.0											
BIC	-1,764.2											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k. Data set limited to a task-specific completion time in the range of 30 to 390 seconds (with the range corresponding to ±2sd from the mean of the completion time distribution).

Table 7: Specification test: basic beta regression mixture model

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.032	***	0.128	-2.781	***	0.085	-2.404	***	0.024	-2.126	***	0.106
`PRIOR`	3.032	***	0.319	4.464	***	0.181	4.808	***	0.005	1.041	***	0.237
(ϕ)_(Intercept)	1.216	***	0.130	3.546	***	0.153	7.272	***	0.326	3.577	***	0.195
(π)_(Intercept)	fixed to 0			-1.488		1.158	-0.661		0.995	-2.872	*	1.610
(π)_`MALE`	fixed to 0			0.465		0.341	0.055		0.350	0.350		0.367
(π)_`AGE`	fixed to 0			-0.021		0.013	-0.022	*	0.013	-0.029	**	0.014
(π)_`EDUCATION`	fixed to 0			0.099	**	0.047	0.065		0.049	0.189	***	0.061
(π)_`NUMERACY`	fixed to 0			1.467	**	0.624	0.238		0.487	1.743		1.117
Membership size[b]	131			420			176			276		
Observations	1,003											
Log-likelihood	-1,003.0											
AIC	-1,936.0											
BIC	-1,764.2											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k.

Table 8: Specification test: basic beta regression mixture model w/o concomitant variables

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.050	***	0.102	-3.016	***	0.088	-2.545	***	0.037	-1.657	***	0.028
PRIOR	2.441	***	0.217	4.692	***	0.180	5.082	***	0.077	0.141	**	0.064
(ϕ)_(Intercept)	1.375	***	0.096	3.991	***	0.189	5.803	***	0.232	6.030	***	0.238
(π)_(Intercept)	0.737	***	0.171	0.573	***	0.173	0.257		0.171	fixed to 0		
Membership size[b]	217			313			293			180		
Observations	1,003											
Log-likelihood	-997.9											
AIC	-1,965.8											
BIC	-1,892.2											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k.

Table 9: Robustness test: basic beta regression mixture model, restricted data set

	Class 1: inconsistent updaters			Class 2: Bayesian learners			Class 3: information refuseniks			Class 4: belief adopters		
	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error	Est.		Std. Error
(Intercept)	-2.032	***	0.128	-2.781	***	0.085	-2.404	***	0.024	-2.126	***	0.106
`PRIOR`	3.032	***	0.319	4.464	***	0.181	4.808	***	0.005	1.041	***	0.237
(ϕ)_(Intercept)	1.216	***	0.130	3.546	***	0.153	7.272	***	0.326	3.577	***	0.195
(π)_(Intercept)	fixed to 0			-1.488		1.158	-0.661		0.995	-2.872	*	1.610
(π)_`MALE`	fixed to 0			0.465		0.341	0.055		0.350	0.350		0.367
(π)_`AGE`	fixed to 0			-0.021		0.013	-0.022	*	0.013	-0.029	**	0.014
(π)_`EDUCATION`	fixed to 0			0.099	**	0.047	0.065		0.049	0.189	***	0.061
(π)_`NUMERACY`	fixed to 0			1.467	**	0.624	0.238		0.487	1.743		1.117
Membership size[b]	131			420			176			276		
Observations	1,003											
Log-likelihood	-1,003.0											
AIC	-1,936.0											
BIC	-1,764.2											

Notes: $^{***}P < 0.01, ^{**}P < 0.05, ^{*}P < 0.1$; [a] mean-centered; [b] membership size determined based on the number of subjects assigned with $\tau_{ik}^{(*)} = \max \tau_{iu}^{(*)} \forall u$ to class k. Data set limited to a task-specific completion time in the range of 30 to 390 seconds (with the range corresponding to ±2sd from the mean of the completion time distribution).

Table 10: Information selection criteria: main model

	Main model			Main model w/o concomitant variables			Main model, restricted data set		
	AIC	BIC	ICL	AIC	BIC	ICL	AIC	BIC	ICL
$K = 1$	-1,526.8	-1,502.2	-1,502.2	-1,526.8	-1,502.3	-1,502.3	-1,488.0	-1,463.7	-1,463.7
$K = 2$	-1,747.6	-1,673.9	-1,659.4	-1,754.7	-1,700.7	-1,686.2	-1,695.1	-1,622.2	-1,606.8
$K = 3$	-1,854.4	-1,731.6	-1,716.4	-1,848.8	-1,765.3	-1,750.1	-1,803.9	-1,682.4	-1,667.0
$K = 4$	-1,936.0	**-1,764.2**	**-1,748.9**	-1,958.1	**-1,845.2**	**-1,830.0**	-1,878.5	**-1,708.4**	**-1,693.1**
$K = 5$	-1,940.4	-1,719.4	-1,703.9	-1,966.7	-1,824.3	-1,808.8	-1,919.9	-1,701.2	-1,685.9
$K = 6$	**-1,962.1**	-1,692.0	-1,676.4	**-1,983.6**	-1,841.2	-1,825.8	**-1,938.7**	-1,671.4	-1,656.1

Notes: AIC = Akaike Information criterion; BIC = Bayesian Information criterion; ICL = Integrated Classification Likelihood criterion (see McLachlan and Peel 2000).

Table 11: Information selection criteria: basic model

	Basic model			Basic model w/o concomitant variables			Basic model, restricted data set		
	AIC	BIC	ICL	AIC	BIC	ICL	AIC	BIC	ICL
$K = 1$	-1,529.8	-1,515.1	-1,515.1	-1,529.8	-1,515.1	-1,515.1	-1,491.8	-1,477.2	-1,477.2
$K = 2$	-1,755.2	-1,701.2	-1,686.7	-1,762.4	-1,728.1	-1,713.5	-1,700.7	-1,647.3	-1,632.8
$K = 3$	-1,884.7	-1,791.4	-1,776.5	-1,893.4	-1,839.3	-1,824.5	-1,833.5	-1,741.2	-1,726.4
$K = 4$	-1,932.0	-1,799.4	-1,784.1	-1,965.8	**-1,892.2**	**-1,877.0**	-1,897.2	**-1,766.0**	**-1,751.0**
$K = 5$	-1,985.0	**-1,813.1**	**-1,797.8**	**-1,972.8**	-1,879.5	-1,864.1	-1,924.8	-1,754.7	-1,739.5
$K = 6$	**-2,000.2**	-1,789.1	-1,773.5	nc	nc	nc	**-1,938.2**	-1,729.3	-1,713.9

Notes: AIC = Akaike Information criterion; BIC = Bayesian Information criterion; ICL = Integrated Classification Likelihood criterion (see McLachlan and Peel 2000); nc = non-convergence after 5,000 iterations.

References

[1] Adda, J. and F. Cornaglia (2006). Taxes, cigarette consumption, and smoking intensity. *American Economic Review 96*(4), 1013–1028.

[2] Ai, C. and E. C. Norton (2003). Interaction terms in logit and probit models. *Economics Letters 80*(1), 123–129.

[3] Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008). Eliciting risk and time preferences. *Econometrica 76*(3), 583–618.

[4] Andersson, H. and P. Lundborg (2007). Perception of own death risk. *Journal of Risk and Uncertainty 34*(1), 67–84.

[5] Arrow, K. J. (1982). Risk perception in psychology and economics. *Economic Inquiry 20*(1), 1–9.

[6] Barron, G. and E. Yechiam (2009). The coexistence of overestimation and underweighting of rare events and the contingent recency effect. *Judgment and Decision Making 4*(6), 447–460.

[7] Bruhin, A., H. Fehr-Duda, and T. Epper (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica 78*(4), 1375–1412.

[8] Cameron, T. A. (2005). Updating subjective risks in the presence of conflicting information: an application to climate change. *Journal of Risk and Uncertainty 30*(1), 63–97.

[9] Cawley, J. and C. Ruhm (2011). The economics of risky health behaviors. *NBER Working Paper 17081.*

[10] Davidson, R. and J. G. MacKinnon (1984). Convenient specification tests for logit and probit models. *Journal of Econometrics 25*(3), 241–262.

[11] Dayton, C. M. and G. B. Macready (1988). Concomitant-variable latent-class models. *Journal of the American Statistical Association 83*(401), 173–178.

[12] Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society B 39*(1), 1–38.

[13] Dickie, M. and S. Gerking (1996). Formation of risk beliefs, joint production and willingness to pay to avoid skin cancer. *Review of Economics and Statistics 78*(3), 451–463.

[14] Eil, D. and J. M. Rao (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics 3*(2), 114–138.

[15] Ferrari, S. and F. Cribari-Neto (2004). Beta regression for modelling rates and proportions. *Journal of Applied Statistics 31*(7), 799–815.

[16] Fischhoff, B. and W. Bruine De Bruin (1999). Fifty–fifty= 50%? *Journal of Behavioral Decision Making 12*(2), 149–163.

[17] Gilbert, D. T. and T. D. Wilson (2007). Prospection: experiencing the future. *Science 317*(5843), 1351–1354.

[18] Harris, A. J. and U. Hahn (2011). Unrealistic optimism about future life events: a cautionary note. *Psychological Review 118*(1), 135–154.

[19] Hogarth, R. M. and H. J. Einhorn (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology 24*(1), 1–55.

[20] Johnson, E. J. and A. Tversky (1983). Affect, generalization, and the perception of risk. *Journal of Personality and Social Psychology 45*(1), 20.

[21] Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review 93*(5), 1449–1475.

[22] Lichtenstein, S., P. Slovic, B. Fischhoff, M. Layman, and B. Combs (1978). Judged frequency of lethal events. *Journal of Experimental Psychology: Human learning and memory 4*(6), 551–578.

[23] Loewenstein, G. F., E. U. Weber, C. K. Hsee, and N. Welch (2001). Risk as feelings. *Psychological Bulletin 127*(2), 267–286.

[24] Lundborg, P. and B. Lindgren (2004). Do they know what they are doing? risk perceptions and smoking behaviour among swedish teenagers. *Journal of Risk and Uncertainty 28*(3), 261–286.

[25] Magat, W. A. and W. K. Viscusi (1992). *Informational approaches to regulation.* Cambridge, MA: MIT Press.

[26] Manski, C. F. (2004). Measuring expectations. *Econometrica 72*(5), 1329–1376.

[27] McFadden, D. (1999). Rationality for economists? *Journal of Risk and Uncertainty 19*(1-3), 73–105.

[28] McFadden, D. (2006). Free markets and fettered consumers. *American Economic Review 96*(1), 3–29.

[29] McLachlan, G. J. and D. Peel (2000). *Finite mixture models.* New York, NY: Wiley.

[30] Mobius, M. M., M. Niederle, P. Niehaus, and T. Rosenblat (2011). Managing self-confidence: Theory and experimental evidence. *NBER Working Paper 17014*.

[31] O'Donoghue, T. and M. Rabin (2001). Choice and procrastination. *Quarterly Journal of Economics 116*(1), 121–160.

[32] Peters, E., D. Västfjäll, P. Slovic, C. Mertz, K. Mazzocco, and S. Dickert (2006). Numeracy and decision making. *Psychological Science 17*(5), 407–413.

[33] Savage, L. (1954). *The Foundations of Statistics.* New York, NY: John Wiley.

[34] Scharff, R. (2012). Economic burden from health losses due to foodborne illness in the united states. *Journal of Food Protection 75*(1), 123–131.

[35] Shogren, J. and T. Stamland (2007). Valuing lives saved from safer food - a cautionary tale revised. *American Journal of Agricultural Economics 89*(5), 1176–1182.

[36] Sloan, F. A., V. K. Smith, and D. H. Taylor (2003). *The smoking puzzle: Information, risk perception, and choice.* Cambridge, MA: Harvard University Press.

[37] Slovic, P. (1999). Trust, emotion, sex, politics, and science: Surveying the risk-assessment battlefield. *Risk Analysis 19*(4), 689–701.

[38] Slovic, P., M. L. Finucane, E. Peters, and D. G. MacGregor (2004). Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis 24*(2), 311–322.

[39] Slovic, P., B. Fischhoff, and S. Lichtenstein (2000). Facts and fears: Understanding perceived risk. In P. Slovic (Ed.), *The Perception of Risk*, pp. 137–154. London: Earthscan.

[40] Slovic, P. and S. Lichtenstein (1971). Comparison of bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance 6*(6), 649 – 744.

[41] Smith, V., D. Taylor, and F. Sloan (2001). Longevity expectations and death: Can people predict their own demise? *American Economic Review 91*(4), 1126–1134.

[42] Smith, V. K. and F. R. Johnson (1988). How do risk perceptions respond to information? the case of radon. *Review of Economics and Statistics 70*(1), 1–8.

[43] Smithson, M., E. C. Merkle, and J. Verkuilen (2011). Beta regression finite mixture models of polarization and priming. *Journal of Educational and Behavioral Statistics 36*(6), 804–831.

[44] Smithson, M. and C. Segale (2009). Partition priming in judgments of imprecise probabilities. *Journal of Statistical Theory and Practice 3*(1), 169–181.

[45] Smithson, M. and J. Verkuilen (2006). A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods 11*(1), 54–71.

[46] Spinnewijn, J. (2013). Insurance and perceptions: how to screen optimists and pessimists. *Economic Journal 123*(569), 606–633.

[47] Vaillant, V., H. de Valk, E. Baron, T. Ancelle, P. Colin, M. C. Delmas, B. Dufour, R. Pouillot, Y. Le Strat, P. Weinbreck, E. Jougla, and J. C. Desenclos (2005). Foodborne infections in france. *Foodborne Pathogens and Disease 2*(3), 221–232.

[48] Viscusi, W. (1989). Prospective reference theory: Toward an explanation of the paradoxes. *Journal of Risk and Uncertainty 2*(3), 235–264.

[49] Viscusi, W. and J. K. Hakes (2008). Risk beliefs and smoking behavior. *Economic Inquiry 46*(1), 45–59.

[50] Viscusi, W. K. (1997). Alarmist decisions with divergent risk information. *Economic Journal 107*(445), 1657–1670.

[51] Viscusi, W. K. (1998). *Rational Risk Policy.* New York, NY: Oxford University Press.

[52] Viscusi, W. K. and C. J. O'Connor (1984). Adaptive responses to chemical labeling: Are workers bayesian decision makers? *American Economic Review 74*(5), 942–956.

[53] Wisdom, J., J. S. Downs, and G. Loewenstein (2010). Promoting healthy choices: Information versus convenience. *American Economic Journal: Applied Economics 2*(2), 164–178.