



HAL
open science

Do People Really Want to Be Informed? Ex-ante Evaluations of Information-Campaign Effectiveness

Romain Espinosa, Jan Stoop

► **To cite this version:**

Romain Espinosa, Jan Stoop. Do People Really Want to Be Informed? Ex-ante Evaluations of Information-Campaign Effectiveness. *Experimental Economics*, 2021, 24 (4), pp.1131-1155. 10.1007/s10683-020-09692-6 . halshs-03097601

HAL Id: halshs-03097601

<https://shs.hal.science/halshs-03097601>

Submitted on 5 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Do People Really Want to Be Informed?

Ex-ante Evaluations of Information-Campaign Effectiveness*

Romain Espinosa

CNRS, CREM - Université de Rennes 1, France

Rennes School of Business, France

Jan Stoop

Erasmus School of Economics, Netherlands

November 13, 2020

Abstract

We develop a method to assess population knowledge about any given topic. We define, and rationalize, types of beliefs that form the ‘knowledge spectrum’. Using a sample of over 7,000 UK residents, we estimate these beliefs with respect to three topics: an animal-based diet, alcohol consumption and immigration. We construct an information-campaign effectiveness index (ICEI) that predicts the success of an information campaign. Information resistance is greatest for animal-based diets, and the ICEI is highest for immigration. We test the predictive power of our ICEI by simulating information campaigns, which produces supportive evidence. Our method can be used by any government or company that wants to explore the success of an information campaign.

Key words: Information campaigns, information resistance, experiment, diet, alcohol, immigration, knowledge spectrum.

JEL Classification: C70, C90, D10, D1.

*We would like to thank Georg Granic, Nicolas Treich, Andrew Clark, the editor and three anonymous referees for useful comments, as well as participants in various seminars (BETA Strasbourg, LESSAC Dijon, and CREM Rennes) and conferences (ASFEE 2019 Toulouse). We are also grateful to ERIM, NWO (Veni grant 016.155.026, and Vidi grant 195.061), ANR (grant ANR-19-CE21-0005-01) for financial support.

1 Introduction

Public policies usually aim to modify deeply-anchored behaviors in the population. Laws and taxes are the government’s most direct way of changing behavior, but they are not always very popular in the eyes of the general public. The Yellow Vests movements in France, Belgium and the UK are examples of how the public opinion has turned against the government (Anderson, 2019). On the other hand, information campaigns are a softer measure that the government can use. The tool is easier to implement, but is by no means cheap. For example, the UK government plans to run 140 information campaigns in 2018-2019 (Griggs, 2018), with annual expenditures being over £300 million (Reynolds, 2014).

The impact of information campaigns is obviously limited when individuals do not want to be informed about the topic. This may be the case when cognitive dissonance is strong (Festinger, 1962; Brehm & Cohen, 1962; Bem, 1967; Elliot & Devine, 1994; Cooper, 2007, e.g.). We propose a method to quantify the relative weight of information resistance, of which cognitive dissonance is a large part, in a particular knowledge domain. More specifically, we propose a question chart that distinguishes various dimensions of the knowledge spectrum. We separate seven types of beliefs, varying from pretending not to have knowledge to wrongfully thinking one has accurate knowledge (see Figure 1, and Appendix A for a rationalization).

After developing the knowledge spectrum, we present an experimental method to disentangle the seven types of beliefs in a sample of over 7,000 UK residents. They are presented with a series of statements and have to report whether each statement is true or false. We exogenously vary if subjects are incentivized for their answers (stated vs. revealed beliefs), and if we provide an information link to a website showing the correct answer. Based on the experimental outcomes, we construct a measure of how effective information campaigns will be, the information-campaign effectiveness index (ICEI). We then use these ICEI’s to predict the rank order of the impact of information campaigns using a new set of experiments.

We apply our method to three topics for which individuals might engage in motivated reasoning, while the government has an interest in having a well-informed population. These topics are animal-based diets, alcohol consumption and immigration. The psychology literature extensively discussed the presence of cognitive dissonance for animal-based diets, and we suspected a potential for resistance to information for alcohol consumption and immigration as well. We summarize our results as follows. First, across topics, we find sharp differences in the presence of information resistance. It is most prevalent in animal-based diets, at roughly 12 percent of our subjects’ knowledge spectrum. This is in line with studies in psychology (Loughnan et al., 2014; Graça et al., 2016; Piazza & Loughnan, 2016). For alcohol, information resistance is close to 0 percent, but rises to almost 3 percent for heavy drinkers. Immigration falls at 1.5 percent.

Second, as a consequence of the first result, we find that the ICEI differs across the topics. Most room for an effect of an information campaign is expected for immigration,

because subjects report low levels of knowledge and low levels of information resistance. These findings are consistent with recent works on information effects on the perception of immigration (Facchini et al., 2016; Grigorieff et al., 2018). Lower levels of improvement are expected for alcohol consumption, because of the combination of good prior knowledge and little information resistance. Low levels of improvement are expected for animal-based diets too. Subjects report a low level of knowledge, but face fierce information resistance.

Third, we test these predictions by simulating information campaigns in an additional set of experiments. The results show that information campaigns do indeed have the greatest impact for immigration. The impact is smaller and similar for diet and alcohol consumption. Last, we present several robustness checks on our measures of information resistance and the ICEI. Among these checks, we show that our results are not driven by (i) respondents exerting extra cognitive effort in the incentivized treatments, (ii) respondents conducting side online searches, (iii) specific questions in our questionnaires, (iv) an experimenter demand effect, (v) an order effect of the control covariates we gather.

The simplicity and effectiveness of our method make it an attractive tool for public authorities, and even private companies, that wish to launch information campaigns. This is the first method that can predict *ex-ante* the potential success of information campaigns. Given the budget constraints that governments face, our method can help to set priorities over competing information campaigns. Moreover, our method provides governments with an idea of the potential opposition that reforms may face: topics with higher information resistance scores are likely to receive more criticism. Our protocol brings therefore more information than what pilot information campaigns would: it does not solely measure the *effectiveness* of a specific information campaign, but assesses the maximum *effectiveness* an information campaign on a specific topic could reach. Finally, this method can be used as a selection device, helping governments choose appropriate arguments to back up policies.

One main component of information resistance is cognitive dissonance. We are not the first to consider these topics in economics. However, we do move this literature in a new direction by using an incentivized method that allows us to measure its extent. Cognitive dissonance was first introduced in economics by Akerlof & Dickens (1982). The concept was used to enrich economic theory, providing ways in which cognitive dissonance at the personal level affects behavior at the aggregate level. Examples can be found regarding issues such as social change (Rabin, 1994; Benabou & Tirole, 2006; Oxoby, 2004), duty-oriented consumption (Nyborg, 2011) and the labor-market consequences of discrimination (Goldsmith et al., 2004).

In empirical and experimental evidence, dissonance is induced exogenously without incentives. Differences in treatments confirm the existence of cognitive dissonance, but remain silent about its magnitude. Contributions in this spirit can be found in the fields of finance (Goetzmann & Peles, 1997; Antoniou et al., 2013; Chang et al., 2016), other-regarding behavior (Konow, 2000; Dickinson & Oxoby, 2011; Matthey & Regner, 2011), and social and

political attitudes (Di Tella et al., 2007; Elinder, 2012). Finally, our work is closely related to recent works in political literature that show that incentivized surveys substantially reduce the gap in stated beliefs between groups of partisans (Bullock et al., 2015; Prior et al., 2015).

In the remainder of the paper, section 2 presents our theoretical background. Section 3 explains the experimental method, and section 4 shows the results. Section 5 tests the predictions of our ICEI estimates, section 6 shows robustness checks and section 7 concludes.

2 Disentangling information resistance

One of the key behaviors our study aims at unraveling is information resistance. Golman et al. (2017) explain that information avoidance can be rational if beliefs are part of the utility function and present several theories that might support it (e.g., risk aversion, anxiety, regret aversion, optimism maintenance, dissonance avoidance). In case an action generates negative externalities, individuals might avoid getting informed or facing the truth whenever their uninformed utility is greater than the expected informed utility.

Information resistance represents a real challenge for policy-makers, as it undermines the effectiveness of information campaigns in two ways. First, some individuals who are informed about the topic might refuse to accept the evidence they were exposed to. As Golman et al. (2017) recall, people might engage in a biased interpretation of information or they might actively seek to forget what they saw. Second, other individuals might actively avoid getting informed on a topic, even when the information is freely accessible (see also Ho et al. (2018)). Information campaigns will be ineffective with these two types of agents, so it would be beneficial for a policy-maker to have an estimate of the extent such agents are present.

To correctly assess the potential of an information campaign, policy-makers need to answer the following questions. What proportion of people knows that the action generates negative externalities? Among them, how many would not admit that it actually generates negative externalities? If they do not know for sure that it generates negative externalities, do they have a clue that it might be the case? If yes, do they actively avoid getting informed on the topic?

We present these questions in the decision tree depicted in Figure 1. The questions allow classifying the answers of the population in a knowledge spectrum with seven ‘beliefs’. The spectrum of knowledge presents seven categories of beliefs that are defined based on (i) the private information people have, (ii) their knowledge about their potential ignorance, (iii) their willingness to get informed, and (iv) their propensity to lie (potentially to themselves) about the information. In Appendix A, we present a simple theory that rationalizes these seven types of belief. Here, we briefly discuss each belief type.

Accepted Knowledge corresponds to the proportion of the population that knows this behavior generates negative externalities and that accepts this fact. On the contrary, *Fake*

Ignorance is associated with the share of individuals who have the knowledge about the negative externalities but who don't report it. These individuals can either report that they don't know whether the action generates negative externalities (*Fake "I don't know"*), or they can report that they believe there are no associated negative externalities (*Fake Incorrect*). These individuals might have various reasons for misreporting their knowledge about the existence of negative externalities.

Of course, individuals can be sincerely ignorant about the externalities they generate. *Unconscious Ignorance* corresponds to the proportion of the population who, sincerely, believe that their behavior does not generate negative externalities. People can, however, know that they have limited information on the topic, and that they are actually ignorant. Among these individuals, some of them actively look to get informed on the topic if they can access to information (*Information-seeking accepted ignorance*). Some conscious ignorant individuals agree to get informed on the topic should accessing the information not be (too) costly or time-consuming (*Information-passive accepted ignorance*). For these individuals, the opportunity cost of getting informed is too large: when they have no incentive to look for information on the topic at stake, they recognize that they are ignorant but do not engage in knowledge acquisition. A third group of agents can, however, actively look to remain uninformed and, thus, avoid information on the topic (*Selective ignorance*).

The aim of this paper is to get a quantitative estimate of information resistance and a measure of how effective an information campaign can be. We thus define two measures of interest. The first is information resistance, which is made of *Selective ignorance*, *Fake "I don't know"* and *Fake Incorrect*. The second is an Information-Campaign Effectiveness Index (ICEI). This index is defined as the sum of beliefs that can be altered by an information campaign, i.e. beliefs associated with *Sincere Ignorance*. By definition, this index takes a value between 0 and 1, with larger values reflecting a greater potential to influence public knowledge. We summarize our two measures as follows:

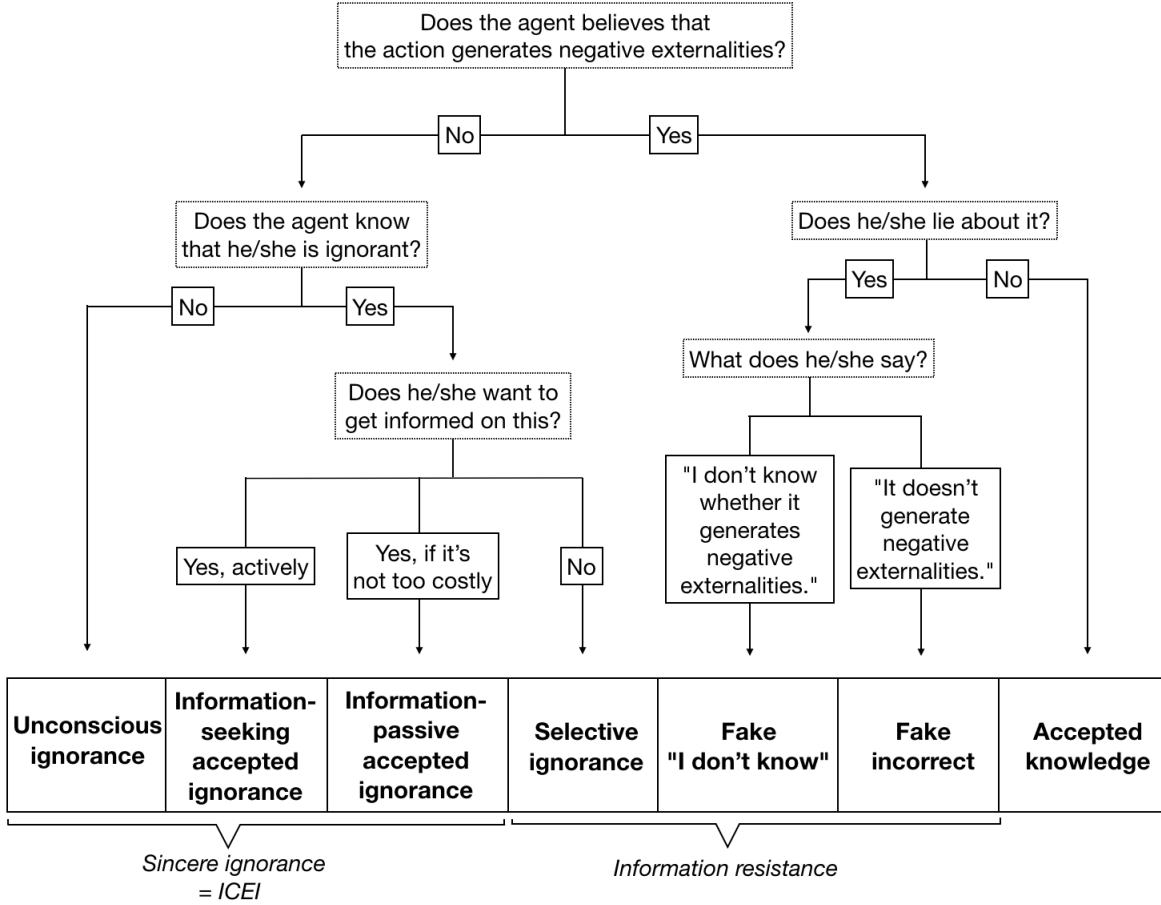
Information Resistance The sum of beliefs categorized as *Fake I Don't Know*, *Fake Incorrect* and *Selective Ignorance*.

Information-Campaign Effectiveness Index (ICEI) The sum of beliefs categorized as *Information-seeking Accepted Ignorance*, *Information-passive Accepted Ignorance* and *Unconscious Ignorance*.

3 An Experimental Method to Identify the Knowledge Spectrum

Section 2 shows how (prior) beliefs can be categorized into seven categories that constitute the knowledge spectrum. Here, we describe the experimental procedure to sort the population's

Figure 1 The decision tree and knowledge spectrum



beliefs into these categories. We first introduce, for a given topic, questionnaires that policy makers could address to the population. We then show how each type of belief is distributed over correct and incorrect answers in the questionnaire. Last, we show how we retrieve these categories using Maximum Likelihood Estimation.

3.1 The different questionnaires and treatments

In order to elicit the elements of the spectrum knowledge, we propose a baseline condition with additional treatment variations. Consider a series of issues for which policy makers want to evaluate the population’s knowledge. We reframe them as a questionnaire consisting of a list of statements that are to be presented to a sample of the population. We consider the case where public authorities can rely on a solid corpus of scientific research or facts to effectively assess the veracity of each statement. The participants are asked to report whether they believe each statement in the questionnaire to be *True* or *False*. We denote by $Pr[C|Q]$

the proportion of correct answers given by the entire sample, for a given questionnaire Q that is subject to treatment variations.

We start from a baseline questionnaire similar to an opinion poll where individuals are asked whether they believe some statements to be true or false. In a naive world, policy makers would use the percentage of incorrect answers to this questionnaire to assess the general ignorance of the population. However, they would overestimate the level of ignorance as they would count pretended ignorance as real ignorance. We propose several treatment variations that aim at disentangling the elements of the knowledge spectrum, and, thus, at better assessing real ignorance. First, by giving participants the possibility to say ‘I don’t know’, we are able to capture the proportion of people who acknowledge their ignorance or those that fake it. Second, by introducing incentives, we are able to force fake ignorants to reveal their knowledge. Third, by giving participants links that contain the necessary information to assess the veracity of the statements, we capture the share of the population willing to get informed.

Table 1 summarizes the three core dimensions of the nine treatments. The first treatment variation concerns the incentive structure: either there is no reward for a correct answer (HYPOTHETICAL) or there are incentives for a correct answer (INCENTIVE). The hypothetical treatments allow for information resistance, as not answering truthfully is costless. However, the incentives make truthful answers increasingly rewarding, and so increase the cost of (self-)deception. The incentives here must be chosen so that they i) are rewarding enough to make people reveal their belief, and ii) provide a ceiling beyond which additional incentives do not increase truthful answers. We choose three levels of incentives: 10, 30 or 50 pence per correct answer (see section 4.1).

The second treatment variation is the option to obtain information about the statement in the questionnaire (for other work on information avoidance and moral issues, see Grossman & Van Der Weele (2017) and Serra-Garcia & Szech (2019)). In the NOLINK treatments, participants are exposed to the questionnaire without hints. In the LINK treatments, however, we provide information in the form of a link to a website that directly addresses the statement at hand. In our experiment, information access does not come at a monetary cost, but incurs opportunity costs in terms of time to obtain the information. The last treatment variation expands the strategy space for respondents. In all of the treatments reported in Table 1, subjects are given two possible answers: true or false. The only exception is treatment H-NL-3. This is a hypothetical treatment with no information link, where subjects have the option to answer ‘I don’t know’.

3.2 The identification strategy

The treatment variations allow identifying the prevalence of the different beliefs presented in the knowledge spectrum (see Figure 1). For each treatment, our theory is able to predict

Table 1 Core design of the treatment

Treatment	Payoff	Information link	Strategy space
H-NL	£0	No	True - False
H-L	£0	Yes	True - False
I-NL10	£0.10	No	True - False
I-NL30	£0.30	No	True - False
I-NL50	£0.50	No	True - False
I-L10	£0.10	Yes	True - False
I-L30	£0.30	Yes	True - False
I-L50	£0.50	Yes	True - False
H-NL-3	£0	No	True - False- I don't Know

H: Hypothetical, I: Incentive, L: Link, NL: No Link. For example, the abbreviation I-NL50 represents a treatment where the incentives are £0.50 per correct answer and no information link is provided in order to assess the truth of the statement.

whether each part of the knowledge spectrum will yield *correct* or *incorrect* answers. Table 2 sets out the parts of the spectrum that appear in each treatment. The rows show the answers that subjects can give, separately for each treatment; the columns show the types of belief in the knowledge spectrum. A column is marked X when a subject's answer can be explained by the column variable. From this, we derive eleven equations to estimate the seven unknown parts of the spectrum.

The first column of Table 2 is associated with *Accepted Knowledge*, which, by definition, only yields correct answers. Therefore, the belief type *Accepted Knowledge* only appears in rows where that represent correct answers. Second, our method describes *Information Seeking Ignorance* as beliefs corresponding to participants (i) who would say 'I don't know' should they have the opportunity to say it and would gamble otherwise, but that (ii) would give the correct answer if they had an access to the information. It is therefore associated with 'I don't know' in H-NL-3, with correct answers in the LINK treatments, and with both correct and incorrect answers in the remaining treatments.

Similarly, *Information Passive Ignorance* is associated with participants who would also like to answer 'I don't know', but would look to get informed if they had access to information only if they are incentivized to do so. This type of belief will therefore yield 'I don't know' answers in H-NL-3, correct answers in the INCENTIVE-LINK treatments, and both correct and incorrect answers in the remaining treatments. The third component of *Sincere Ignorance* is *Unconscious Ignorance*. Our method describes it as individuals who are sincerely convinced of the incorrect answer. Thus, incentives and links are ineffective as participants do not manipulate their beliefs nor do they suspect to be misinformed.

Fourth, information resistance is made of three types of beliefs. Our method describes *Fake 'I don't know's* as beliefs that yield participants to answer 'I don't know' whenever they

Table 2 Summary of the identification strategy

Treatment	Accepted Knowledge	Sincere Ignorance			Information Resistance		
		Info. Seeking	Info. Passive	Unconscious Ignorance	Fake IDK	Fake Incorrect	Selective Ignorance
$P(\text{Correct} \text{H-NL-3})$	x						
$P(\text{IDK} \text{H-NL-3})$		x	x		x		
$P(\text{Incorrect} \text{H-NL-3})$				x		x	x
$P(\text{Correct} \text{H-NL})$	x	x	x		x		
$P(\text{Incorrect} \text{H-NL})$		x	x	x	x	x	x
$P(\text{Correct} \text{I-NL})$	x	x	x		x	x	
$P(\text{Incorrect} \text{I-NL})$		x	x	x			x
$P(\text{Correct} \text{H-L})$	x	x	x		x		
$P(\text{Incorrect} \text{H-L})$			x	x	x	x	x
$P(\text{Correct} \text{I-L})$	x	x	x		x	x	x
$P(\text{Incorrect} \text{I-L})$				x			

$P(\cdot|\cdot)$ represents the probability of an answer, conditional on the treatment. The answers can be correct, incorrect, or ‘I don’t know’. The treatment abbreviations are H: hypothetical, I: Incentivized, L: information link available and NL: no information link available; the strategy space is of 3 answers.

can, to give correct answers when they are incentivized, and to gamble in the other cases. It is therefore associated with ‘I don’t know’ in H-NL-3, with correct answers in the INCENTIVE treatments, and with both correct and incorrect answers in the remaining treatments. Similarly, beliefs associated with *Fake Incorrect* lead participants to answer incorrectly as long as they are not incentivized, and to give correct answers when they are. It yields correct answers in INCENTIVE treatments, and incorrect answers in the other treatments.

Last, participants subject to *Selective Ignorance* have an incorrect view of the issue at stake, but they refuse to get informed if they are not incentivized to do so. This category of beliefs leads therefore to incorrect answers in all conditions, except in INCENTIVE-LINK treatments.

3.3 Maximum-Likelihood Estimation

Table 2 sets out a set of conditions that allows estimating the belief categories by Maximum-Likelihood Estimation (MLE). Our MLE program computes the likelihood of each equation given the individual data, assuming the normal distribution of the error term ($y \sim N(\mu, \sigma)$). Each mean μ is expressed as a combination of types of beliefs, as given by the conditions in Table 2 (see Appendix E). We constrain the set of beliefs and strategies to take values between 0 and 1 (inclusive), using the *L-BFGS-B* algorithm in R (package `optim`). We estimate μ and σ assuming that a proportion s of *Sincere Ignorance* and a proportion s' of *Fake IDK* will report the correct state of the world when they cannot reply *I don’t know*. We weigh observations (normalized to H-NL) to ensure that each condition is given the same weight in

the MLE.

We iterate the estimation process following a bootstrapping procedure with $B = 10,000$ repetitions with replacement. For each treatment containing N_t observations, we randomly draw N_t observations from the sample with replacement. We average the estimates to obtain average estimated beliefs over the B repetitions. From the empirical distribution of the estimated coefficients, we calculate the proportion of estimates that are equal to zero. We then carry out one-sided tests by rejecting the null hypothesis ($H_0 : \beta = 0$) for the alternative hypothesis ($H_1 : \beta > 0$) if the proportion of estimated parameters equal to zero is below α . Note that one-sided bootstrap tests in the case of restricted MLE exhibit the correct asymptotic null-rejection rates (Andrews, 2000).

4 Applications: animal-based diets, alcohol and immigration

To showcase our methodology, we apply it to three topics that we believe are of interest to governments: animal-based diets, alcohol and immigration. A growing body of scientific evidence calls for less consumption of animal-based products, mostly for environmental, health and animal-welfare reasons (Espinosa, 2019, e.g.). Reducing meat consumption can sharply reduce greenhouse gas emissions (Tilman & Clark, 2014; Aleksandrowicz et al., 2016; Springmann et al., 2016) and deforestation (Erb et al., 2016), and can improve health and reduce health-care costs (Berners-Lee et al., 2012; Hu et al., 2000; Song et al., 2004; Schulze et al., 2005; Aune et al., 2009; Sinha et al., 2009; Elwood et al., 2010; Pan et al., 2011, 2012; Micha et al., 2012; Larsson & Orsini, 2013; Orlich et al., 2013).

One of the main advantages of our methodology is that it measures the degree of information resistance. In our setting, we mostly consider that the government wants to raise awareness about some individual behaviors that generate negative externalities on society. In these cases, we believe that cognitive dissonance is the most important driver of information resistance, as individuals want to avoid dissonance between what they would like to do and what their morals urge them to do. Therefore, we present below an overview of the literature on cognitive dissonance regarding the three topics that we consider.

Numerous contributions in psychology have found that the consumption of animal-based products, especially meat, is subject to cognitive dissonance (Loughnan et al., 2014). As individuals are attached to the consumption of animal-based products, they underestimate the negative externalities of their consumption on the environment, health, and/or animal welfare, even when they are provided with information on the topic. Psychological research has shown that individuals engage in motivated reasoning regarding meat consumption (Loughnan et al., 2014; Graça et al., 2016; Piazza & Loughnan, 2016) and for genetically-modified food (Fernbach et al., 2019), and has explored the types of arguments used to justify their consumption (Joy, 2011; Rothgerber, 2014; Piazza et al., 2015; Caviola et al., 2019).

The negative consequences of alcohol use have been widely documented. In general, country-level alcohol use rises with income, but then starts to fall after a certain income threshold (Holmes & Anderson, 2017). Alcohol use leads to significant health damage (WHO, 2018) and severe economic damage (Bouchery et al., 2011). One effective way of reducing alcohol consumption is to raise its price (Cook & Moore, 2002). Of more interest in the current paper’s context, governments continually develop new public-information strategies on the negative consequences of alcohol use (Morris, 2019, e.g.). Psychologists have studied the link between alcohol use and information resistance. The results are not as strong as those for meat consumption, but there is greater information resistance among heavy drinkers.

Some papers support the idea that information resistance might play a role for those who consume moderate levels of alcohol. McNally et al. (2005) conclude that reducing cognitive dissonance has some effect on reducing drinking behaviour. Mäkelä (1997) shows that motivated reasoning might take place: cognitive dissonance can be alleviated by thinking that others drink more. However, other work suggests that cognitive dissonance might only play a limited role in alcohol consumption. For instance Steele et al. (1981) show that dissonance arousal had little effect on the act of drinking alcohol. Overall, the benefits of lying to oneself are lower for alcohol consumption, as most of the downside effects of consumption (i.e. the health impacts) fall on the individual.

Last, we look at information resistance regarding immigration as it has been showed that there is a large discrepancy between beliefs and numbers on this topic (Alesina et al., 2018). In 2015, at the peak of the refugee crisis in the EU, 1.82 million refugees entered the EU (Quarterly, 2016). The crisis led to social unrest in a number of countries, mostly in the form of local protests. There is mixed evidence, however, on its influence in the political sphere. Gehrsitz & Ungerer (2017) show that there was no increase in votes for the main anti-immigrant party in Germany after an influx of immigrants. Steinmayr (2016) estimate that voting for the main anti-immigrant party did increase in Austria, but fell when locals had more contact with immigrants.

On the other hand, Sekeris & Vasilakis (2016) find that xenophobia did rise when refugees landed in Greece, and that the motives to vote for anti-immigrant parties were xenophobic. There is some evidence on the effectiveness of information campaigns. Facchini et al. (2016) and Grigorieff et al. (2018) show that the support for immigration by nationals can be largely affected by information campaigns. Related to our research, Freddi (2017) is the only paper we are aware of that tests a form of cognitive dissonance regarding immigration. Whereas the paper does not try to estimate the magnitude of information resistance, it reports that Swedes search less often for news articles on refugees on the internet, the closer they live to a refugee centre.

4.1 Experimental population

All treatments are carried out on the platform Prolific. This is an online platform similar to Amazon MTurk where people can subscribe to finish tasks for payment. Payments consist of a fixed hourly wage plus potential incentives. Without loss of generality for our theory, this platform offers two advantages. First, the platform allows testing our methodology on a large sample of subjects. Second, Prolific offers a fine balance between paying subjects too little and too much. Because experiments are short, it is customary that payments are in the range of a couple of pounds.

Prolific offers a sample responsive to small incentives. In the INCENTIVE treatments, we reward participants from £0.10 to £0.50 per correct answer. This range of incentives has shown to overcome partisan biases in other works (Prior et al., 2015; Bullock et al., 2015, e.g.). Our stakes make it plausible that information resistance is felt during the experiment, but that choices are not affected by subjects anticipating that they will suffer from it outside of the experiment. Note that the representativeness of the subjects' population should be a concern for any work aiming to draw policy recommendations. In this paper, however, we merely aim to show a proof-of-principle of the method.

Between one fourth and one fifth of the participants in our sessions are students (see table C.1). This is more than the general population, which has a share of 3.5% (source: <https://www.universitiesuk.ac.uk>). About two thirds of the participants were female, which is more than the roughly 50% females that live in the UK (source: <https://www.statista.com>), and the average age of our subjects lies between 33 and 38 years old, which is close to the median age of 40 (source: <https://www.statista.com>). Half of the respondents had a full-time job and one fourth had a part-time job. Its combination comes close to the UK's employment rate of 76.3% (source: <https://www.ons.gov.uk>). Job-seeking respondents constitute about 10% to 15% of the sample. The samples are relatively similar across treatments.

4.2 Experimental design and procedure

Following the design presented in Table 1, subjects were presented with six questions about either an animal-based diet, alcohol consumption or immigration. All the questions can be found in Appendix B. The questions were asked in the same order for all participants, and no subject participated in more than one treatment. At the end of each survey, we asked subjects about either their dietary habits, alcohol intake or political preferences. The flat rate payment in all of the sessions was £0.20, translating into an hourly wage of £7.12. In addition, the subjects in the treatment with incentives earned an average of £0.97.

All experiments were carried out using the Prolific platform, between October 2018 and January 2019. To obtain a homogenous subject pool, we recruited subjects who satisfied all of the following criteria: UK nationality, current country of residence is the UK, and first language is English. We used a generic recruitment text to invite subjects to the experiment.

We informed them about a questionnaire that would take roughly two minutes to complete, with a possibility of earning a bonus. Therefore, in the hypothetical treatments only, we added an incentivized question independent from the rest of the experiment at the end.

4.3 Main Results

Table 3 shows OLS results of the reduced form regressions. We regress the percentage of correct answers for each of the three topics on a set of dummy variables. Link is a dummy for there being an information link and IDK reflects the treatments with an ‘I don’t know’ answer option (H-NL-3). I_x refers to the incentives for a correct answer: 10, 30 or 50 pence. The coefficients in columns with no controls are the mean differences across treatments (we provide the associated descriptive statistics in Table C3 of Appendix C). In addition to the demographics presented above, the controls ask about diet (ranging from ‘vegan’ to ‘daily meat eater’), alcohol consumption (‘once per month’ to ‘once per day’) and position on a political scale (‘Far Left’ to ‘Far Right’).¹

Table 3 Regression of the percentage of correct answers on the different treatments

	Percentage of correct answers					
	Diet		Alcohol		Immigration	
Link	0.0497*** (0.0131)	0.0532*** (0.0126)	0.0444*** (0.0103)	0.0426*** (0.0104)	0.152*** (0.0120)	0.149*** (0.0121)
I10	0.0645*** (0.0184)	0.0530*** (0.0176)	-0.00457 (0.0147)	-0.00192 (0.0151)	-0.00537 (0.0171)	-0.0118 (0.0173)
I30	0.0446** (0.0186)	0.0400** (0.0180)	-0.00765 (0.0144)	-0.00577 (0.0147)	-0.000353 (0.0171)	-0.00572 (0.0173)
I50	0.0543*** (0.0186)	0.0443** (0.0179)	-0.0177 (0.0144)	-0.0161 (0.0149)	-0.00616 (0.0173)	-0.00933 (0.0173)
IDK	-0.137*** (0.0230)	-0.152*** (0.0209)	-0.122*** (0.0201)	-0.122*** (0.0202)	-0.213*** (0.0203)	-0.221*** (0.0206)
Constant	0.498*** (0.0145)	0.753*** (0.0513)	0.589*** (0.0113)	0.587*** (0.0743)	0.511*** (0.0134)	0.322*** (0.0386)
Controls	No	Yes	No	Yes	No	Yes
Observations	1,795	1,795	1,800	1,800	1,800	1,700
R-squared	0.071	0.173	0.049	0.068	0.190	0.235

Robust standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels. Controls include: age, occupation, gender, student status, and dummy variables for missing demographics. We further control for diet in Diet, alcohol consumption in Alcohol, and political self-placement in Immigration.

Table 3 shows no significant difference across the three incentive levels for all three topics (the associated p-values are all over 0.28). This is reassuring, as it suggests a ceiling effect:

¹We found no statistical differences across treatments within topics for the control questions (Kruskal-Wallis equality-of-populations rank test; diet: $p = 0.321$, alcohol: $p = 0.802$, political self-placement: $p = 0.259$).

offering higher incentives does not make subjects give better answers. Moreover, incentives only have a significant effect for animal-based diets. This hints that information resistance plays a significant role here, but less for the other two topics. The constant in the no-controls specifications is largest for alcohol: subjects have the greatest baseline knowledge on this topic. Finally, the coefficient of Link is largest for immigration. Providing information links produces an improvement in the score that is about three times that in the two other domains.

We next show the MLE results of the knowledge spectrum in Figure 1. The regression results are reported in Table 4. First, information resistance differs across topics, with that for diet being the highest at 12.4% of the knowledge spectrum; there is no information resistance for alcohol consumption or immigration. Second, accepted knowledge also differs across topics, with alcohol taking the lead. Third, the degree to which subjects are willing to seek information is the largest for immigration. All of these findings explain the differences we observe in the ICEI. We will use this index as an input for a direct test of an information campaign, which will be presented in the next section.

Table 4 MLE estimates of the three topics

	Diet		Alcohol		Immigration	
	Entire sample	Bootstrap	Entire sample	Bootstrap	Entire sample	Bootstrap
Accepted Knowledge	0.315	0.315*** (0.016)	0.434	0.434*** (0.015)	0.302	0.300*** (0.015)
Information Seeking	0.134	0.119** (0.051)	0.192	0.183*** (0.025)	0.357	0.348*** (0.021)
Information Passive	0.027	0.050 (0.066)	0	0.005 (0.018)	0	0.003 (0.012)
Unconscious Ignorance	0.390	0.390*** (0.010)	0.373	0.373*** (0.007)	0.342	0.341*** (0.010)
Fake IDK	0.132	0.124*** (0.004)	0	0.004 (0.012)	0	0.004 (0.009)
Fake Incorrect	0	0 (0)	0	0 (0)	0	0.002 (0.006)
Selective Ignorance	0	0 (0)	0	0 (0)	0	0.002 (0.006)
Information Resistance	0.132	0.124	0	0.004	0	0.008
ICEI	0.550	0.558	0.565	0.560	0.699	0.692
N	1,795		1,800		1,800	

Standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels.

5 Testing the Effectiveness of Information Campaigns

5.1 Campaigns across topics

In the above section, we elicited the beliefs of a sample of the UK population about three topics: animal-based diets, immigration, and alcohol consumption. We estimated Information Campaign Effectiveness Index (ICEI) scores, which correspond to the proportion of sincere

incorrect beliefs in the population. We obtained figures of 55.8% for animal-based diets, 56.0% for alcohol consumption and 69.2% for immigration. Our theory now predicts that information campaigns should have a greater effect for immigration than for animal-based diets and alcohol consumption. Had we not taken information resistance into account, the associated ICEI score for animal-based diets would have been almost equivalent to what we observe for immigration.

In what follows, we put these predictions to the test. To do so, we simulate online information campaigns and evaluate the associated change in reported knowledge. We create an infographic for each topic that contains six items, each of which provided the necessary information to correctly answer the six questions in the associated questionnaire. The infographics corresponding to animal-based diets, immigration and alcohol consumption appear in Appendix D. The order of the information presented in the infographics maps the order of the questions in the questionnaire. All infographics are presented in the same order, with the same layout and in the same color.

We define a new treatment, H-NL-Info, where participants are introduced to the infographics before facing the hypothetical no-link questionnaire. This new treatment aims to reproduce real-life information campaigns, where participants see adverts on TV, on the street or on their computer. Participants are not incentivized to memorize the information, such that information-resistant participants can disregard the campaign and can go directly to the questionnaire. We ran this additional treatment in January 2019, with 200 participants for each topic. Participants to this additional treatment had similar demographics to those of the above treatments (see Table C1).² The results of the new treatment are presented in Table 5. Our conjectures, based on Table 4, are confirmed.

Table 5 Impact of simulated information campaigns

	Diet	Alcohol	Immigration
ICEI	0.558	0.560	0.692
H-NL (% correct)	49.9 %	58.5 %	51.9 %
H-NL-Info (% correct)	69.3 %	77.9 %	82.4 %
Campaign Impact	+19.4 pp ($p < 0.01$)	+19.4 pp ($p < 0.01$)	+ 30.5 pp ($p < 0.01$)

The percentage of correct answers in the hypothetical questionnaire increases by 19.4 percentage points for the diet and alcohol questionnaires, while that for immigration rises by 30.5 percentage points. Using OLS, we estimate a diff-in-diff coefficient to assess the statistical difference across treatments. We find no difference in the information campaign effect between diet and alcohol ($p = 0.99$). We do find a stronger impact for immigration than for diet ($p < 0.01$), and for immigration versus alcohol consumption ($p < 0.01$). These

²We found no statistical differences between H-NL and H-NL-Info within topics for the control questions (Mann-Whitney U test; diet: $p=0.39$, alcohol: $p=0.35$, political self-placement: $p=0.50$).

results confirm and validate our ICEI: we observe similar impacts of information campaigns for topics with similar ICEI scores, and a larger effect for topics with higher ICEI scores.

5.2 Predictive power of ICEI and Information Resistance

The above results show that the ICEI score helps to predict on which topic information campaigns will be the most effective in changing beliefs. An important question is whether our method improves the prediction about the effectiveness of an information campaign, as compared to more naive measures. For comparison, we compare our estimates to two naive measures: the share of incorrect answers in H-NL, and the number of clicks in H-L. To do so, we compute for our 18 questions (6 for diet, 6 for alcohol, 6 for immigration) the impact of the infographics and estimate their ICEI and Information Resistance scores. Pooling the 18 questions, we observe a very high and significant correlation between the ICEI score and the actual impact of the infographics ($\hat{\rho} = 0.766, p < 0.001$).

To analyze the gain in predictive power of our method, we regress the impact of the infographics on the share of correct answers on several variables (Table 6). First, we start with the baseline scenario where the policy maker would only consider the share of incorrect answers in a H-NL survey. The higher the share of incorrect answers (*IncorrectHNL*), the more policy makers would expect an information campaign to be useful. We observe indeed a positive relationship between the share of incorrect answers in HNL and the impact of the infographics. The predictive power of *IncorrectHNL* is given by the R^2 score, which is equal to 51.5%. Second, we consider the case where a policy-maker would consider the number of clicks between HNL and HL to get a prediction of the impact of her campaign. The second column shows that the prediction power decreases to 32.3%.

Third, we regress the impact of the infographics on the ICEI score (column 2). We see that the ICEI significantly correlates with the infographics' impact and that the predictive power increases up to 58.6%. Fourth, we adjust our ICEI score. Note, the impact of the infographics is computed on the share of correct answers in H-NL. However, information-seeking and information-passive participants gamble at this stage (as they cannot answer "I don't know") and some of them already give the correct answer. So, we correct the ICEI score by subtracting these individuals using the MLE estimates. The third column shows that this increases the prediction power to 68.1%.

Last, it is possible that our estimates of information resistance only capture a lower bound (e.g., our incentives are too small to deter all psychological mechanisms of information resistance). However, if our estimates of information resistance sufficiently well correlate with the true levels of information resistance, controlling for our Information Resistance estimates should improve the prediction level. Column 4 shows indeed that this is the case: adding the Information Resistance estimate improves the prediction power to 78.7% (+10 percentage points). The adjusted R^2 also increases by almost 10 percentage points, which shows the

merits of adding this control.

Table 6 Regression of the impact of infographics

	(1)	(2)	(3)	(4)	(5)
IncorrectHNL	0.5446*** (0.1321)				
ShareClickHL		0.5033** (0.182)			
ICEI			0.5830*** (0.1225)		
Adjusted ICEI				0.6699*** (0.1146)	0.6274*** (0.0979)
Information Resistance					-0.4188*** (0.1535)
N	18	18	18	18	18
R ²	51.5%	32.3%	58.6%	68.1%	78.7%
Adjusted R ²	48.5%	28.0%	56.0%	66.1%	75.9%

Standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels.

6 Robustness checks

In this section, we propose additional tests to investigate the robustness of the results presented in section 4.3.

First of all, our main identification strategy for information resistance is to compare the answers in the hypothetical to those in the incentivized treatments. However, this comparison might be confounded by two factors. First, economists have long argued that incentives make subjects exert more cognitive effort (Camerer & Hogarth, 1999; Zimmermann, 2019, e.g.). The increase in accuracy could thus result from an increase in cognitive effort in recalling the information at stake. In addition, incentives can cause our subjects to search the internet for the correct answer (see Grewenig et al. (2019) for evidence of such kind of behavior). In this case, the increase in accuracy would reflect a higher frequency of online searches to answer our surveys. These two confounds may cause the increase in accuracy to stem from an increase in effort instead of honesty.

Evidence suggests, however, that our results are not driven by an increase in cognitive effort or online searches effort. To see this, note that subjects take 62.54 seconds to complete a no-link treatment and 135.26 seconds to complete a link treatment (Mann-Whitney U test, $N_1 = 2,405$, $N_2 = 2,396$, $p < 0.01$). So, when we provide subjects the shortest route to the correct answer, answering one question takes on average 12.12 seconds longer.

Second, pooling all treatments, we find that subjects take 95.03 seconds to complete an hypothetical treatment, and 100.10 seconds to complete an incentivized treatment (Mann-

Whitney U test, $N_1 = 1,201, N_2 = 3,600, p = 0.07$). The 0.85 second per question that subjects take longer in the hypothetical treatment cannot therefore reasonably account for internet searches, nor can it represent intensive efforts to better consider the question in the incentivized treatment.

We provide further evidence in Table 7 below, by regressing the percentage of correct answers on treatment dummies and controlling for the total response time (this variable is measured in minutes). As can be seen, adding response time as a control variable does not appreciably change the effect of incentivizing participants. This supports the fact that the increase in accuracy results from increased honesty and does not result from increased (cognitive) efforts. We summarize our finding here:

Robustness check 1 Incentivizing subjects in our treatments does not cause them to exert more (cognitive) effort. Rather, subjects answer more honestly.

Table 7 Regression of the percentage of correct answers on the different treatments controlling for total response time

	Percentage of correct answers					
	Diet		Alcohol		Immigration	
Link	0.0496*** (0.0130)	0.0274** (0.0133)	0.0445*** (0.0104)	0.0198* (0.0109)	0.152*** (0.0119)	0.0833*** (0.0121)
I10/30/50	0.0545*** (0.0150)	0.0496*** (0.0149)	-0.00994 (0.0120)	-0.0107 (0.0118)	-0.00396 (0.0137)	-0.0111 (0.0130)
IDK	-0.137*** (0.0234)	-0.140*** (0.0232)	-0.122*** (0.0191)	-0.121*** (0.0189)	-0.213*** (0.0214)	-0.222*** (0.0203)
Time		0.0222*** (0.00341)		0.0207*** (0.00306)		0.0456*** (0.00317)
Constant	0.498*** (0.0146)	0.477*** (0.0148)	0.589*** (0.0116)	0.569*** (0.0118)	0.511*** (0.0133)	0.472*** (0.0130)
Observations	1,795	1,795	1,800	1,800	1,800	1,795
R-squared	0.071	0.092	0.048	0.072	0.190	0.274

Robust standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels. Note that the three treatments with incentives are merged into a single category I10/30/50.

Second, we take an alternative approach to test whether differences in accuracy are due to information resistance. We investigate how incentivizing participants for a survey where we expect no information resistance impacts the accuracy of their answers. To do so, we design a treatment with questions about mundane topics, such as the number of persons on the first shuttle to the moon, where answering them should not be affected by information resistance (see table B4 in Appendix B for the list of question).

We invited 201 new UK subjects from the platform Prolific, using the same selection criteria as we did before (see section 4.2). Half of the subjects were in the hypothetical treatment

and the other half earned £0.10 per correct answer. Our theory predicts no difference in scores here, as there is no reward to (self-)deception. A Mann-Whitney test does not reject the null hypothesis that scores are equal across the two groups ($N_1 = 100, N_2 = 101, \bar{x}_1 = 0.64, \bar{x}_2 = 0.63, p = 0.832$).

This result further validates the part of our method that relies on incentivizing subjects to overcome information resistance: if incentives would cause subjects to exert more effort to recall information, we would have found that in this treatment where cognitive dissonance should not be present. Likewise, in our core treatments, we find different levels of information resistance across topics. If incentives would trigger behaviors other than cognitive dissonance, we would not have found such differences. In sum:

Robustness check 2 Incentives have no effects for topics that ex-ante are free of information resistance.

Third, we look at the possibility that certain design features of our experiment had a large effect on the outcomes. For example, it is possible that our estimates of the knowledge spectrum are driven by an outlier question that presses the score in a particular direction. Omitting one question each time, we ran all main analyses of Table 4 again to see how our measures of information resistance and the ICEI are influenced.

We find that Information Resistance is stable for animal-based diets, ranging from 10.1% to 15.7% (vs. 12.4% for the entire set of questions). For alcohol and immigration, information resistance remains zero percent. The ICEI scores are stable too, with a range from 0.526-0.587 for diets, from 0.534-0.595 for alcohol, and from 0.685-0.724 for immigration.

Robustness check 3 The measurement of information resistance and the ICEI is not driven by any particular question that subjects answered.

Fourth, although we did our best to frame our questions neutrally and objectively, subjects may suspect that the researchers have a bias regarding the topics in question. An experimenter demand effect may appear after a few questions, once the participants form beliefs about what the experimenter expects. If so, we would expect to see different information resistance and ICEI scores in the first half as compared to the second half of the questionnaire.

To investigate this issue, we re-ran the analysis of Table 4 using the first three questions only. Our estimates of information resistance and the ICEI are virtually the same if we consider the first three questions, or all questions. We do observe a small, but not meaningful, drop of information resistance for animal based diets (12.4 percent for the entire sample versus 9.1 percent for the first three questions). Furthermore, the ICEI for immigration is slightly greater when we consider all questions (0.699) instead of the first three questions only (0.636). We summarize this finding here:

Robustness check 4 Our results are not driven by an experimenter demand effect.

Fifth, another important design feature is when to ask subjects about their relevant background characteristics. For example, if dietary habits are asked at the start of the questionnaire, subjects may start to form beliefs about the aims of the experiment or they might be influenced by recalling them about their social identity. This may change their answers to the questionnaire. Alternatively, participants who are asked ex-post might report differently their consumption habits or political self-placement as they have been influenced by the questionnaire.

We explored this idea for the Diet treatment in a pilot study before running our main set of treatments.³ Here, we randomly asked subjects at the start, or at the end, about their dietary habits. We invited 800 subjects who we randomized into a 2×4 design. The questions were hypothetical, or paid £0.10 per correct answer; provided a link with information or not; or asked subjects what diet they followed at the start or at the end of the questionnaire. Like the main treatments, the diet question was measured on a six point scale, ranging from ‘vegan’ to ‘daily meat eater’.

We ran an OLS regression similar to that reported in Table 3, but adding a dummy indicating the diet question was asked before or after the six statements (see table C4 in Appendix). The constant, the dummy for an information link, and the dummy for providing incentives are all similar to those reported in Table 3. Furthermore, the coefficient of the order of question on diet habit is insignificant. The results remain unaffected when we include the background controls. Therefore, we can rule out any priming effects or the like that can be a result of the order in which background variables are asked. Note, given the absence of an order effect, we decided not to run this check for the treatments on immigration and alcohol, which we ran after the diet treatment. For immigration and alcohol, we only asked background characteristics at the end. We summarize our finding here:

Robustness check 5 Asking subjects about their diet before or after the main questionnaire does not influence responses.

Finally, we test whether people who are expected to suffer more from cognitive dissonance, actually do so. We proceed in two ways. First, we re-estimate the variables in Table 4 using only the subjects that are daily meat eaters, report to drink at least 10 units of alcohol per week, and have a value of 6 for the political scale variable (for the Diet, Alcohol, and Immigration questionnaires respectively). Note that this amount is self-reported, and we cannot check its reliability. For each group, we observe higher level of information resistance than for the entire sample, even though it is not statistically different from zero (see Table C5 in Appendix). The ICEI scores are similar to the estimates of our entire sample.

³Note that this questionnaire included an ‘I don’t know’ answer option.

Second, we used the possibility to pre-screen participants in Prolific to run two new treatments on alcohol consumption: H-NL and I-NL. We use the same selection criteria as described in section 4.2, but we add a restriction on alcohol use: self-reported drinking of at least 10 units of alcohol per week. Because of the limited number of available subjects, we invited 100 heavy drinkers to H-NL and 100 heavy drinkers to I10-NL. We then ran an OLS regression of correct answers on the incentivized treatment dummy variable ($p = 0.524$). Because of this insignificance, information resistance must be insignificant too. Note, we did not run these additional sessions for the Diet and Immigration treatments as we were not able to pre-select daily-meat eaters and extreme-right voters on the platform. We summarize our final robustness check here:

Robustness check 6 Self-reported ‘daily meat eaters’, ‘heavy drinkers’ and ‘heavy right-wing voters’ do not generate significantly higher levels of information resistance.

7 Discussion and Conclusion

Information campaigns are one of the favorite tools of governments to bring about behavioral change in the population. Compared to taxation or hard-law regulation, information campaigns have greater social acceptability, as they leave citizens with the same decision space. The success of such campaigns does, however, rely on the assumption that people are willing to become informed about topics in which the government is interested. Information resistance postulates on the contrary that some individuals might reject the information they receive when accepting it would force them to change their behaviors.

This paper addresses the consequences of information resistance for public authorities. To provide policy-makers with an ex-ante evaluation of the potential impact of information campaigns, we propose a question chart that sorts beliefs into seven categories. In Appendix A, we show that economic theory can easily rationalize these types of beliefs, and we propose an easily-implementable experimental method that allows estimating information resistance. We also construct an Information Campaign Effectiveness Index, which corresponds to the share of incorrect beliefs that individuals are willing to revise.

We apply our method to three topics: (animal-based) diets, alcohol consumption and immigration. Our experiments show a level of information resistance of 12.4 percent for animal-based diets. Regarding the ex-ante effectiveness of information campaigns, we found a high potential for the topic of immigration. Smaller and similar potential effects are found for alcohol consumption and animal-based diets. We simulated three information campaigns to test our predictions, and we verified a significantly stronger campaign effect for immigration than for animal-based diets and alcohol.

The main methodological development of this paper is to propose to policy-makers, NGOs and even private companies an efficient way to develop and assess the effectiveness of informa-

tion campaigns. At very low monetary and time costs, these institutions can easily predict the effectiveness of specific pre-defined arguments. Given the contribution to the ICEI score, they will be able to select the most arguments with the largest impact. Considering information resistance, they will also be able to detect the potential polarization in the population that a campaign may generate. For instance, policy-makers or environmental NGOs could seek to find the best set of arguments to convince citizens to decrease meat consumption, i.e., those that offer the highest ICEI score and the weakest information resistance. Decomposing the data by type of argument allows to identify the furthest-reaching information campaign. As far as our sample is concerned, environmental arguments are good topics to communicate on as they have a large ICEI score with a low level of information resistance (see Table 8 below). On the contrary, health arguments will have a more limited impact as the population is more knowledgeable on the topic and is relatively more resistant to information.

Our work is a first step in economics that seeks to quantify cognitive dissonance. However, future works could build on it to explore several directions. First, we interpret the treatment variations as the result of cognitive dissonance: when it becomes costly to lie, participants reveal the information they have or start to get informed on the topic. While we have excluded the most salient alternative explanations (demand effect, undetected online searches, cognitive effort), future works could test whether alternative forces could drive the results. Second, we consider here "Information Passive Accepted Ignorance" as sincerely ignorant individuals whose opportunity costs of getting informed is too large to actively look for information on the topic. However, the opportunity cost of accessing the information could be strategically used by willful ignorant participants for not getting informed on the topic (Ehrich & Irwin (2005); Grossman & Van Der Weele (2017)). In this regard, our results are likely to capture a lower bound estimation of the Information Resistance. Future works could seek to disentangle this issue. Third, we did not consider here the role of the type of sources provided to the participants. It might be that participants might attribute a lower or greater credibility to governmental, NGO or journalistic sources. Policy makers could use our method to select the best set of sources, i.e., those that maximize the ICEI score.

Table 8 Belief spectrum for the three types of arguments related to diet

	Animal Welfare	Environment	Health
Accepted Knowledge	0.294	0.227	0.438
Information Seeking	0.156	0.031	0.071
Information Passive	0	0.273	0.015
Unconscious Ignorance	0.404	0.392	0.376
Fake IDK	0.145	0.073	0.1
Fake Incorrect	0	0	0
Selective Ignorance	0	0	0
Information Resistance	0.145	0.073	0.10
ICEI	0.559	0.696	0.462

Appendix A: A theoretical model of cognitive dissonance for the spectrum of knowledge

In this section we develop a theoretical model based on cognitive dissonance theory that would account for the beliefs reported in the spectrum of knowledge (see Figure 1).

General framework. We consider a situation where agents can consume a private good that may generate negative externalities for society. In the good state of the world ($s = 0$), private consumption c does not generate negative externalities, while it does in the bad state ($s = 1$). We assume here that the true state of the world is $s = 1$. In the following, we take our inspiration from the models of memory management in Bénabou & Tirole (2002) and Hestermann et al. (2018) and assume that individuals are made up of two agents (Self 0 and Self 1).⁴ The true state of the world is not directly observable by either Self 0 or Self 1.

Self 0 receives a number of signals over time about the state of the world and constructs a belief, δ_0 , about the probability that the current state of the world is bad ($\delta_0 = Pr[s = 1]$). Self 0 can fully or partially reveal its knowledge to Self 1, but hiding knowledge is costly (that means, there is a *self-deception cost* *The self-deception costs might vary across individuals, across topics, and can regroup several psychological factors relative to information acquisition such as the preference for the truth.*). This process mimics motivated reasoning in cognitive-dissonance theory, where individuals try to underestimate the negative externalities of their behavior to justify their consumption of the private good. The resulting belief of Self 1 is denoted β .

The timing of the game is as follows: (i) Based on its belief δ_0 , Self 0 decides on the information to transmits to Self 1, which generates beliefs β , and (ii) based on its beliefs β , Self 1 chooses consumption of the potentially-harmful private good. In the following, we consider an equilibrium for this game where both Self 0 and Self 1 maximize their utility, and where Self 1 naively trusts the information provided by Self 0.⁵ The equilibrium is calculated by backward induction: Self 1 maximizes its utility given the beliefs transmitted by Self 0, and Self 0 chooses the optimal level of motivated reasoning anticipating Self 1's choice.

The left-hand side of Figure A1 illustrates Self 1's utility depending on its consumption of the private good and its belief β . This depicts the utility functions of Self 1 for four types of beliefs: $\beta = 1 > \beta_1 > \beta_2 > \beta = 0$. These curves reflect a number of fundamental assumptions

⁴In Bénabou & Tirole (2002), individuals have two distinct time horizons, which we model here as two Selves, as in Hestermann et al. (2018).

⁵We consider Self 1 here as a naive non-Bayesian agent. There are two reasons for doing so. First, our objective is to develop a model that accounts for non-Bayesian behaviors (i.e. cognitive dissonance). Specifying a model with a Bayesian agent to account for non-Bayesian behaviors is as debatable as specifying a naive agent. Second, Hestermann et al. (2018) show that modeling Self 1 as a Bayesian agent still leaves room for Self 0 to manipulate the beliefs of Self 1 at some cost. As the two models (naive agent and Bayesian agents) yield the same type of results (manipulating beliefs is effective, but at some cost), we choose simplicity here and model Self 1 as a naive agent.

about agent preferences. First, the concavity of the utility curves reflects the tradeoff between decreasing marginal utility and the fixed price of consuming the private good ($\frac{\partial^2 U(c|\beta)}{\partial c^2} < 0$). Second, all utilities are equal (and normalized to 0) when the individual consumes nothing, as zero consumption generates no externalities whatever the state of the world ($\forall \beta, U(0|\beta) = 0$).

Third, for each given level of consumption of the private good, utility is larger when the individual believes that the bad state of the world (i.e. generating negative externalities) is less likely. More formally, we have that $\forall (\beta_1, \beta_2) \in (0, 1)^2$, such that $\beta_1 > \beta_2, \forall c \geq 0 : U(c|\beta_1) < U(c|\beta_2)$. We then define the convex function $V_1(\beta)$ as the maximum utility reached by Self 1 for a given belief β ($V_1(\beta) = \max_c \hat{A} U(c|\beta)$). At $\delta = 0$ and $\delta = 1$, the state of the world is certain, and Self 1 derives utility $V_1(0)$ and $V_1(1)$ respectively. See the left-hand side of Figure A1.

The right hand-side of Figure A1 shows how Self 0 can manipulate Self 1's beliefs. Following Hestermann et al. (2018), we assume that Self 0's utility is equal to Self 1's utility minus a cognitive cost that rises with the degree of belief manipulation ($V_0(\delta_0, \beta) = V_1(\beta) - \phi(\delta_0, \beta)$, $\frac{\partial \phi(\delta_0, \beta)}{\partial \beta} \leq 0$ for $\beta < \delta_0$). The graph starts with the utility of agent Self 0 if Self 1's beliefs are equal to Self 0's beliefs ($\beta = \delta_0$) (i.e. no manipulation). The dashed line represents the utility of Self 0 given his original belief δ_0 . When Self 0 does not manipulate beliefs (i.e. $\beta = \delta_0$) and, therefore does not bear any cognitive costs, the utilities of Self 1 and Self 0 are the same.


If Self 0 does manipulate beliefs, Self 1 will underestimate the probability of the bad state of the world, which will increase his utility (V_1 increases), and, by definition, Self 0's utility as well. However, Self 0 will also incur a growing cognitive cost associated with this belief manipulation ($\phi(\delta, \beta)$ increases). The optimal manipulated belief $\beta^*(\delta_0)$ corresponds to the argument that maximizes the utility of Self 0 ($\beta^*(\delta_0) = \operatorname{argmax}_\beta V_0(\beta, \delta_0)$). We characterize below the types of beliefs in the equilibrium where Self 0's strategy is $\beta^*(\delta_0)$ and Self 1's strategy is $c^* = \operatorname{argmax}_c \hat{A} U(c|\beta^*(\delta_0))$.

We now consider the case where individuals are asked by an interviewer to report the state of the world. We assume a general strategy where Self 1 will report $s = 1$ if his belief is above a threshold $\bar{\delta}$, and $s = 0$ if his belief is below a threshold $\underline{\delta}$. We allow individuals to report 'I don't know' when their beliefs take values between $\underline{\delta}$ and $\bar{\delta}$.

Definition: Accepted knowledge in the population corresponds to the share of individuals who correctly and spontaneously report the state of the world ($s = 1$). This is equal to: $Pr[\delta_0 \geq \bar{\delta}; \beta^*(\delta_0) \geq \bar{\delta}]$, which simplifies to $Pr[\beta^*(\delta_0) \geq \bar{\delta}]$.

The proportion of *Accepted knowledge* in society depends on the distribution of original beliefs δ_0 and the cost of self-deception. Note that in the definition, the first condition is always fulfilled when the second condition is met, since $\beta^*(\delta_0) \leq \delta_0$. First, the more the distribution is concentrated close to $\delta = 1$ (left-skewed), i.e. the more the population truly believes that $s = 1$ is likely, the greater the share of *accepted knowledge*. Second, a higher cost

Figure A1 Utility curves and motivated reasoning



of self-deception is associated with a greater share of *accepted knowledge*. Higher cognitive self-deception costs reduce the incentives to lie, reducing the proportion of individuals whose true beliefs are above $\bar{\delta}$ but whose manipulated beliefs are below $\bar{\delta}$.

Definition : *Pretended ignorance* in the population corresponds to the proportion of individuals whose original beliefs produce $s = 1$ but whose manipulated beliefs lead them to report either ‘I don’t know’ or $s = 0$. The proportion of fake ‘I don’t knows’ is $Pr[\delta_0 \geq \bar{\delta}; \beta^*(\delta_0) \in [\underline{\delta}, \bar{\delta}]]$, and the proportion of fake incorrect reporting ($s = 0$) is $Pr[\delta_0 \geq \bar{\delta}; \beta^*(\delta_0) < \underline{\delta}]$. The total amount of *pretended ignorance* is $Pr[\delta_0 \geq \bar{\delta}_0; \beta^*(\delta_0) < \bar{\delta}]$.

The amount of *pretended ignorance* in society is positively correlated with the amount of evidence supporting $s = 1$ and negatively associated with self-deception costs. For the former, pretended ignorance can only occur when original beliefs support $s = 1$. Individuals do not need motivated reasoning if most of the evidence supports $s = 0$, as their original belief δ_0 is already below $\bar{\delta}$. This reduces the likelihood that the first condition is met. For the latter, higher costs make self-deception less attractive for individuals. Individuals whose original beliefs are above $\bar{\delta}$ are less likely to manipulate their beliefs such that $\beta^*(\delta_0) < \bar{\delta}$. Higher self-deception costs therefore reduce the likelihood that the second condition is met.

Definition : *Accepted ignorance* in the population corresponds to the proportion of individuals whose original and manipulated beliefs lead them to report ‘I don’t know’. This is equal to $Pr[\beta^*(\delta_0) \in [\underline{\delta}, \bar{\delta}]; \delta_0 \in [\underline{\delta}, \bar{\delta}]]$.

The share of *accepted ignorance* also depends on the distribution of original beliefs and self-deception costs. On the one hand, relatively-undocumented questions, or topics with contradictory evidence, will produce a greater mass of individuals between the two thresholds

$\underline{\delta}$ and $\bar{\delta}$. This will increase the probability that the first condition is met. On the other hand, greater self-deception costs reduce the incentives to engage in motivated reasoning, and so the distortion between δ_0 and $\beta^*(\delta_0)$. This will increase the probability of the second condition, given that the first condition is fulfilled.

Imagine now that we provide participants with the possibility to access information revealing the correct state of the world. We assume that information access has an opportunity cost k . This opportunity cost might vary across individuals, depending on various psychological factors such as curiosity or entertainment. Individuals who acknowledge that they are ignorant might be willing to access this information if the benefits of doing so outweigh the associated costs. Access has two opposite effects on individual utility. The first effect is that utility may rise if the signal reveals a good state of the world ($\forall \delta_0 \in (0, 1) : V_0(0) \geq V_0(\beta^*(\delta_0))$).


Alternatively individuals may discover a bad state of the world, which will force them to reduce their level of consumption and engage in cognitive dissonance ($\forall \delta_0 \in (0, 1) : V_0(\beta^*(1)) \leq V_0(\beta^*(\delta_0))$). Self 0 expects to discover a bad state of the world with probability δ_0 and a good state of the world with probability $1 - \delta_0$. The expected utility of information access is then $(1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1)) - k$. Individuals who acknowledge that they are ignorant will access the information if $(1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1)) - k > V_0(\beta^*(\delta_0))$. This is summarized in Figure A2, which illustrates the case where individuals are willing to become informed.

Definition : Population *Information-seeking accepted ignorance* is the share of individuals whose original and manipulated beliefs lead them to report ‘I don’t know’, but who would actively access information about the true state of the world. This is $Pr[\beta^*(\delta_0) \in [\underline{\delta}, \bar{\delta}]; \delta_0 \in [\underline{\delta}, \bar{\delta}]; (1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1)) - k \geq V_0(\beta^*(\delta_0))]$.

Definition : Population *Information-passive accepted ignorance* is the share of individuals whose original and manipulated beliefs lead them to report ‘I don’t know’, and who would not actively access information about the true state of the world. This is $Pr[\beta^*(\delta_0) \in [\underline{\delta}, \bar{\delta}]; \delta_0 \in [\underline{\delta}, \bar{\delta}]; (1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1)) - k < V_0(\beta^*(\delta_0))]$.

Both types of *accepted ignorance* depend on the distribution of original beliefs and self-deception costs, as discussed above. The decomposition of *accepted ignorance* into *information-seeking* and *information-passive* categories depends on the costs of information access and the convexity of the maximum utility function. Higher information costs will naturally discourage access. The condition $(1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1)) - k > V_0(\beta^*(\delta_0))$ is less likely to hold for larger k . In addition, the convexity of $V(\cdot)$ will reduce the second term of the equation $V_0(\beta^*(\delta_0))$, so that revealing the state of the world will become more attractive and $(1 - \delta_0)V_0(0) + \delta_0 V_0(\beta^*(1))$ will be more likely to be above $V_0(\beta^*(\delta_0))$. Overall, *information-seeking accepted ignorance* is expected to rise with the convexity of $V(\cdot)$.

Figure A2 Information Access: information-seeking accepted ignorance



Definition : *Selective ignorance* corresponds to the share of individuals whose manipulated beliefs lead them to report $s = 0$, and who, therefore, do not want to become informed, but would have accessed informed in the absence of cognitive dissonance. The proportion of *selective ignorance* in society is $Pr[\beta^*(\delta_0) < \underline{\delta}; \delta_0 \in [\underline{\delta}, \bar{\delta}]; (1 - \delta_0)V_0(0) + \delta_0 V_0(1) - k \geq V_0(\delta_0)]$.

The population frequency of *selective ignorance* depends on the distribution of original beliefs, the cost of self-deception, the cost of information, and the convexity of the maximum utility function. This frequency first rises with the proportion of individuals whose original beliefs lie between $\underline{\delta}$ and $\bar{\delta}$. The less well-known is the topic for the general population, the more likely is *selective ignorance*. Second, we expect a negative relationship between *selective ignorance* and the costs of self-deception: these latter reduce the distortion between δ_0 and $\beta^*(\delta_0)$, which makes $\beta^*(\delta_0) < \underline{\delta}$ less likely to hold.

Third, it is straightforward to see that higher costs of information access reduce the willingness to be informed. This reduces the likelihood that the last condition holds, leading to a negative relationship between information costs and *selective ignorance*. Last, we noted above that greater convexity in the utility function makes information more attractive. We therefore expect a positive relationship between *selective ignorance* and the convexity of the maximum utility function.

Definition : *Unconscious ignorance* refers to the population share whose original and manipulated beliefs support $s = 0$. This proportion of individuals is $Pr[\beta^*(\delta_0) < \underline{\delta}; \delta_0 < \underline{\delta}]$, which simplifies to $Pr[\delta_0 < \underline{\delta}]$.

The proportion of *unconscious ignorance* in society depends only on the distribution of original beliefs: the more evidence there is for $s = 0$, the more the distribution of beliefs will

be concentrated towards $\delta = 0$, leading to higher shares of *unconscious ignorance*. Note that *unconscious ignorance* does not depend on self-deception costs, as individuals in this category would also support $s = 0$ if they did not engage in motivated reasoning. Finally, we consider here that individuals who are unconsciously ignorant are sufficiently sure that $s = 0$ that the opportunity costs of information access outweigh the associated benefits.

Appendix B: The questions in the questionnaire

Table B1 Statements in the online questionnaire: Meat consumption in the UK

Order	Question (as statements)	Correct answer
1	Eating meat and fish is essential for the healthy growth of children. link: https://www.ncbi.nlm.nih.gov/pubmed/19562864	False
2	Processed meat is a carcinogen (i.e., food that causes cancer). link: http://www.who.int/features/qa/cancer-red-meat/en/	True
3	Should everyone erase meat from their diets, greenhouse gas emissions would decrease by 10% at most. link: https://www.drawdown.org/solutions/food/plant-rich-diet	False
4	The vast majority of soy production (which is the most important factor for deforestation) is used to produce vegan food items (like tofu). link: http://wwf.panda.org/wwf_news/?247051	False
5	For a cow to give milk independently, the most important requirement is that she eats grass. link: https://www.theguardian.com/commentisfree/2017/mar/30/dairy-scary-public-farming-calves-pens-alternatives	False
6	Unlike caged production, organic eggs do not involve the killing of young male chicks (inappropriate for egg-laying). link: https://www.independent.co.uk/life-style/food-and-drink/hatched-discarded-gassed-what-happens-to-male-chicks-in-the-uk-10088509.html	False

Table B2 Statements in the online questionnaire: Alcohol consumption in the UK

Order	Question (as statements)	Correct answer
1	Frequent drinking is more common among low income earners than among high income earners (earning $\hat{\text{£}}40,000$ and above annually). www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/drugusealcoholandsmoking/bulletins/opinionsandlifestylesurveyadultdrinkinghabitsingreatbritain/2005to2016	False
2	Alcohol causes extra risk of heart attacks in men only when they drink more than 10 pints of average strength beer a week. link: https://www.nhs.uk/news/lifestyle-and-exercise/people-who-drink-above-uk-alcohol-guidelines-lose-one-two-years-life/	False
3	Alcohol dependency occurs more than twice as much in the UK as in the rest of Europe. link: http://www.who.int/substance_abuse/publications/global_alcohol_report/profiles/gbr.pdf	True
4	Alcohol is a carcinogen (i.e. consumption of it causes cancer). link: https://www.cancer.gov/about-cancer/causes-prevention/risk/alcohol/alcohol-fact-sheet#q2	True
5	People who use alcohol before age 15 are as likely to become alcohol dependent as adults who begin drinking at age 21. link: https://www.cdc.gov/alcohol/faqs.htm#young	False
6	Less than 5% of the traffic accidents in the UK involve at least one driver being over the drink-drive limit. link: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/402698/rrcgb-drink-drive-2013-prov.pdf	False

Table B3 Statements in the online questionnaire: Immigration

Order	Question (as statements)	Correct answer
1	The UK's visa entry requirements for skilled workers have been lowered in the past ten years. link: https://ichef.bbci.co.uk/news/624/cpsprodpb/7E95/production/_101150423_chart-imm-entry_v2-nc.png	False
2	Foreign-born women in the UK on average give birth to about twice as many children as UK-born women. link: https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/livebirths/bulletins/parentscountryofbirthenglandandwales/2017	False
3	About 8% of the workers in the UK's construction sector are non-UK nationals. link: https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/internationalmigration/articles/migrantlabourforcewithinthecreationindustry/august2018	True
4	Asylum seekers in the UK receive at least £10 per day for food, sanitation and clothing. link: https://www.unhcr.org/asylum-in-the-uk.html	False
5	Most of the Non-EU immigrants that come to the UK, come to look for a job. link: https://ichef.bbci.co.uk/news/624/cpsprodpb/002B/production/_101134000_reasonsforcoming-nc.png	False
6	The UK forms the second-largest receiver of asylum applications in the EU (with Germany being first). link: https://www.bbc.com/news/world-europe-44660699	False

Table B4 Statements in the online questionnaire: Control questions that should trigger no information resistance

Order	Question (as statements)	Correct answer
1	The first shuttle to the moon had three persons on board.	True
2	This soccer stadium is located in France. [Picture of Allianz Arena]	False
3	Toyota belongs to the top 10 largest companies in the world (ranked by revenue).	True
4	In 1990, the UK had between 56 and 58 million inhabitants.	True
5	In 2011, the UK's main source of generating electricity was coal.	False
6	The company Google receives more income from licences than from other sources (e.g. advertisements).	False

Appendix C: Statistics table

Table C1 Demographics

	Main treatments			Infographics		
	Diet	Alcohol	Immigration	Diet	Alcohol	Immigration
Student						
Yes (%)	15	17.8	17.8	23.4	23.4	25.8
No (%)	82.3	78.2	79.1	74.6	74.6	72.2
Missing (%)	2.7	4.1	3.1	2	2	2
Gender						
Male (%)	35.2	29.4	35.9	31.8	25.4	29.3
Female (%)	64	68.3	62.4	67.2	73.1	70.2
Missing (%)	0.9	2.2	1.7	1	1.5	0.5
Age						
Mean	38.1	38	36.5	32.6	34.5	33.5
Std Dev.	(12.1)	(13.2)	(12.3)	(10.4)	(11.2)	
Missing (%)	2.6	5.1	3.8	2.5	2.5	2.5
Employment						
Start next month (%)	0.8	0.9	1.3	0	1.5	0
Full-time (%)	49.5	47.1	49.6	48.8	48.8	48
Unemployed - not seeking (%)	15.4	14.7	11.4	13.4	12.4	11.1
Other (%)	4.1	4.7	5.2	4.5	2	4
Part-time (%)	23.2	21.8	22.6	25.9	23.4	27.3
Unemployed - seeking (%)	4.5	5.9	6.6	6	10	6.6
Missing (%)	2.6	4.9	3.3	1.5	2	3
Observations	1,795	1,800	1,800	201	201	198

Table C2 Number of participants in each treatment

Treatment	Diet	Alcohol	Immigration
<i>Main</i>			
H-NL	199	200	197
H-NL-3	201	188	200
H-L	200	206	199
I10-NL	199	200	201
I10-L	198	201	200
I30-NL	198	208	200
I30-L	199	199	200
I50-NL	200	205	198
I50-L	201	193	200
<i>Pilot</i>			
H-NL	191	-	-
H-L	211	-	-
I10-NL	217	-	-
I10-L	181	-	-
<i>Heavy Drinkers</i>			
H-NL	-	200	-
I10-NL	-	200	-
<i>Infographics</i>			
H-NL	201	201	198
<i>No cognitive dissonance</i>			
H-NL		201	

Table C3 Summary statistics

	Diet		Alcohol		Immigration	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Percentage of correct answers						
H-NL	.499	.253	.585	.199	.519	.218
H-L	.547	.274	.638	.214	.669	.266
I10-NL	.572	.235	.579	.212	.496	.216
I10-L	.603	.276	.634	.209	.668	.255
I30-NL	.54	.25	.576	.205	.516	.208
I30-L	.595	.273	.631	.204	.658	.260
I50-NL	.544	.259	.586	.215	.516	.209
I50-L	.61	.267	.600	.187	.647	.270
H-NL-3	.361	.253	.467	.227	.298	.216
Percentage of 'I don't know' answers						
H-NL-3	.334	.241	.224	.224	.352	.275

Table C4 Regression of the percentage of correct answers for the pilote sessions on the order of the screen.

	Percentage of correct answers	
	Diet	
Link	0.0724*** (0.0192)	0.0673*** (0.0184)
I10	0.0826*** (0.0191)	0.0779*** (0.0183)
Ask diet first	-0.00316 (0.0191)	0.000221 (0.0184)
Constant	0.337*** (0.0188)	0.667*** (0.0464)
Diet	No	Yes
Observations	800	800
R-squared	0.037	0.117

Robust standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels.

Table C5 MLE estimates of the knowledge spectrum for daily meat eaters (Diet), heavy drinkers (Alcohol) and extreme-right voters (Immigration)

	Diet	Alcohol	Immigration
Accepted Knowledge	0.257*** (0.018)	0.431*** (0.036)	0.214*** (0.029)
Information Seeking	0.095 (0.060)	0.209*** (0.045)	0.310*** (0.059)
Information Passive	0.061 (0.075)	0.001 (0.008)	0.001 (0.01)
Unconscious Ignorance	0.442*** (0.014)	0.329*** (0.021)	0.425*** (0.024)
Fake IDK	0.144*** (0.053)	0.015 (0.031)	0.046 (0.052)
Fake Incorrect	0 (0)	0.015 (0.02)	0.002 (0.008)
Selective Ignorance	0 (0)	0.001 (0.006)	0 (0.001)
Information Resistance	0.144*** (0.053)	0.031 (0.034)	0.048 (0.052)
ICEI	0.598*** (0.052)	0.539*** (0.046)	0.736*** (0.057)
<i>N</i>	977	214	289

Bootstrap estimates. Standard errors appear in parentheses. ***/**/*: significant at the 1%/5%/10% levels.

Appendix D: The infographics

Figure D1 Infographic for the diet treatment

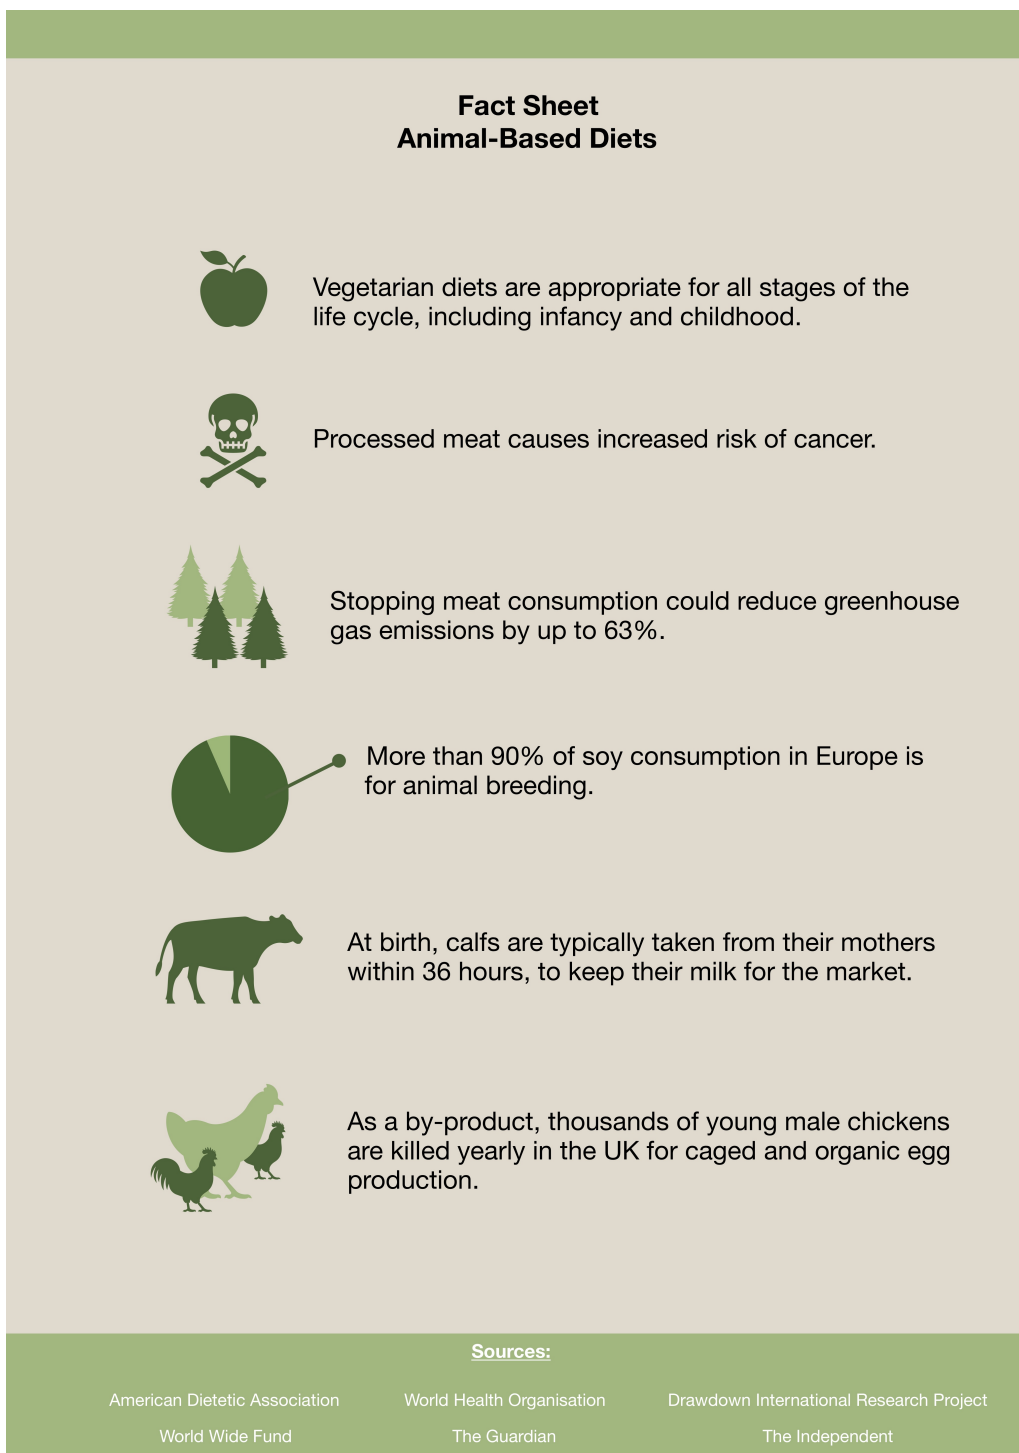
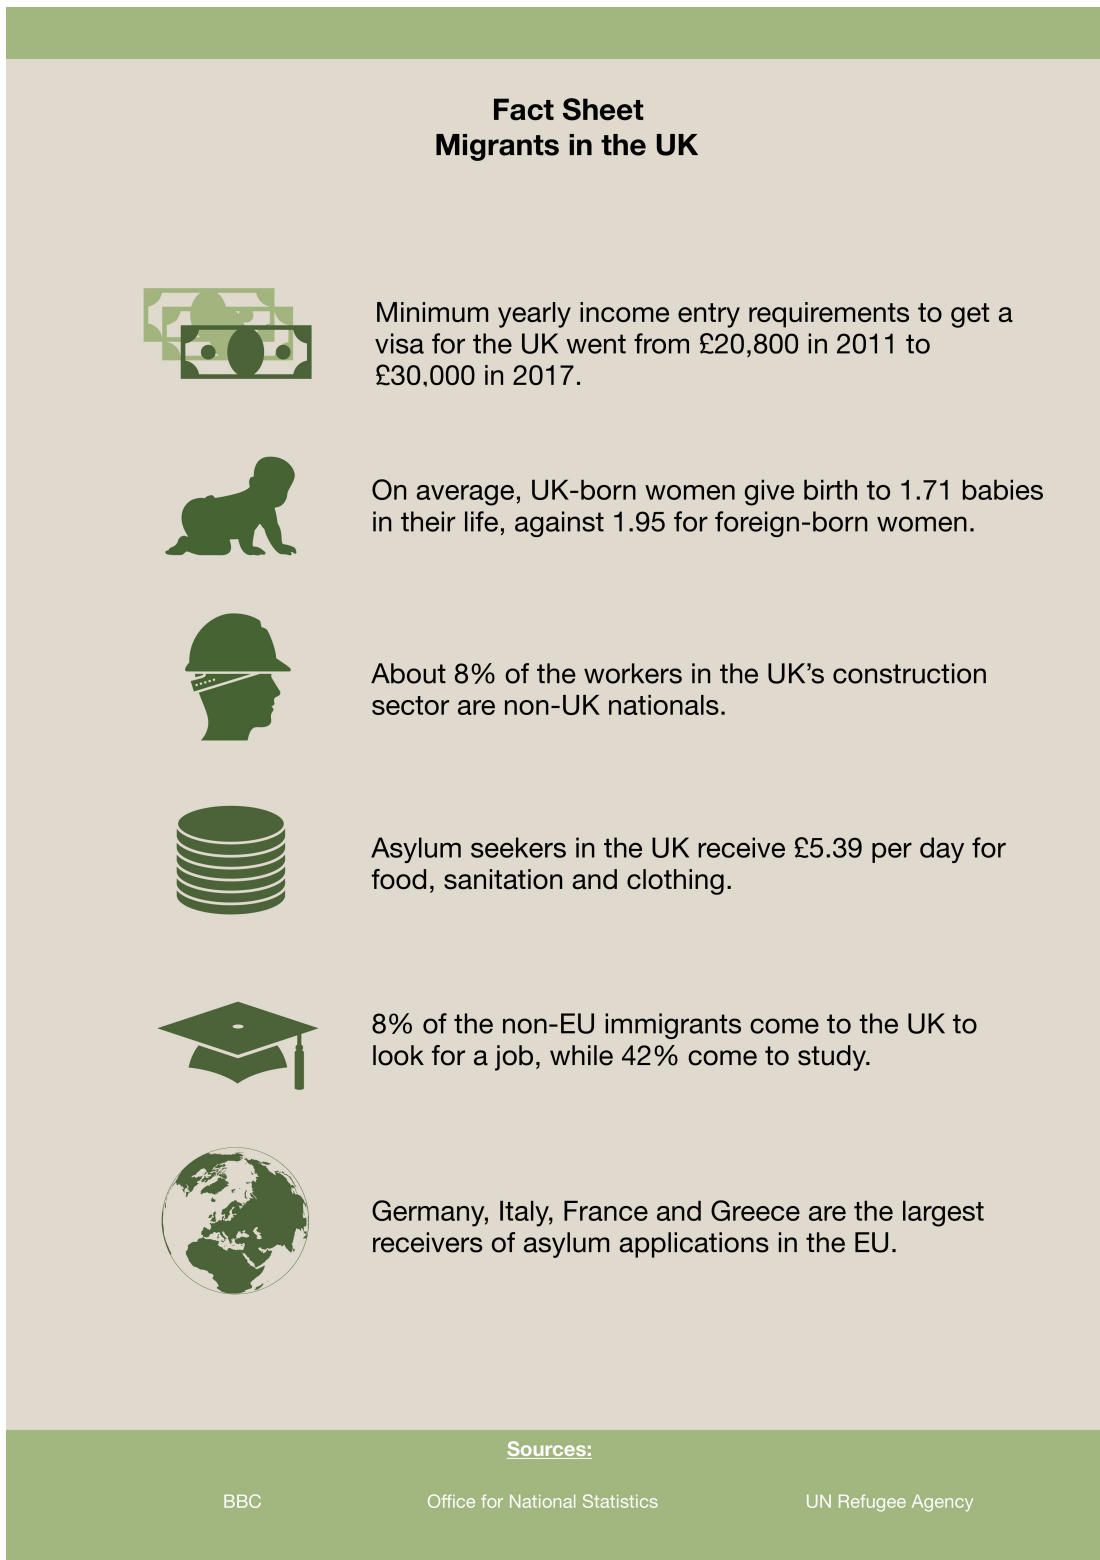


Figure D2 Infographic for the alcohol treatment



Figure D3 Infographic for the immigrant treatment



Appendix E: Maximum Likelihood Estimation

Table 2 summarizes how each type of belief is expressed as a correct, incorrect, or ‘I don’t know’ answer. Let us note ϕ_k the proportion of each type of belief in the knowledge spectrum. Table 2 can be rewritten in forms of equations:

$$P(\text{Correct}|\text{H-NL-3}) = \phi_{AK}$$

$$P(\text{IDK}|\text{H-NL-3}) = \phi_{IS} + \phi_{IP} + \phi_{FIDK}$$

$$P(\text{Incorrect}|\text{H-NL-3}) = \phi_{UI} + \phi_{FI} + \phi_{SI}$$

$$P(\text{Correct}|\text{H-NL}) = \phi_{AK} + s(\phi_{IS} + \phi_{IP}) + (1 - s')\phi_{FIDK} + \phi_{FI}$$

$$P(\text{Incorrect}|\text{H-NL}) = (1 - s)(\phi_{IS} + \phi_{IP}) + (1 - s')\phi_{FIDK} + \phi_{FI} + \phi_{UI} + \phi_{SI}$$

$$P(\text{Correct}|\text{I-NL}) = \phi_{AK} + s(\phi_{IS} + \phi_{IP}) + \phi_{FIDK} + \phi_{FI}$$

$$P(\text{Incorrect}|\text{I-NL}) = (1 - s)(\phi_{IS} + \phi_{IP}) + \phi_{UI} + \phi_{SI}$$

$$P(\text{Correct}|\text{H-L}) = \phi_{AK} + \phi_{IS} + s\phi_{IP} + s'\phi_{FIDK}$$

$$P(\text{Incorrect}|\text{H-L}) = (1 - s)\phi_{IP} + (1 - s')\phi_{FIDK} + \phi_{FI} + \phi_{UI} + \phi_{SI}$$

$$P(\text{Correct}|\text{I-L}) = \phi_{AK} + \phi_{IS} + \phi_{IP} + \phi_{UI} + \phi_{FIDK} + \phi_{FI} + \phi_{SI}$$

$$P(\text{Incorrect}|\text{I-L}) = \phi_{UI}$$

Let us note y_i the proportion of correct, incorrect, or ‘I don’t know’ answers given to the questionnaire by individual i . For each treatment j , we assume that the answers are normally distributed with mean μ_j ($y \sim N(\mu_j, \sigma)$). The contribution to the likelihood for the proportion of correct answers of individual i in treatment j is $f(\frac{y_i - \mu_j}{\sigma_j})$ with $f(\cdot)$ the pdf of a normal distribution.

References

- Akerlof, G. A. & Dickens, W. T. (1982). The economic consequences of cognitive dissonance. American Economic Review, 72(3), 307–319.
- Aleksandrowicz, L., Green, R., Joy, E. J., Smith, P., & Haines, A. (2016). The impacts of dietary change on greenhouse gas emissions, land use, water use, and health: a systematic review. PloS one, 11(11), e0165797.
- Alesina, A., Miano, A., & Stantcheva, S. (2018). Immigration and redistribution. Technical report, National Bureau of Economic Research.
- Anderson, K. B. (2019). France’s yellow vests: A self-mobilised mass movement with insurrectionist overtones. Green Left Weekly, (1207), 1.
- Andrews, D. W. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. Econometrica, 68(2), 399–405.
- Antoniou, C., Doukas, J. A., & Subrahmanyam, A. (2013). Cognitive dissonance, sentiment, and momentum. Journal of Financial and Quantitative Analysis, 48(1), 245–275.
- Aune, D., Ursin, G., & Veierød, M. (2009). Meat consumption and the risk of type 2 diabetes: a systematic review and meta-analysis of cohort studies. Diabetologia, 52(11), 2277–2287.
- Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. Psychological Review, 74(3), 183.
- Bénabou, R. & Tirole, J. (2002). Self-confidence and personal motivation. Quarterly Journal of Economics, 117(3), 871–915.
- Benabou, R. & Tirole, J. (2006). Belief in a just world and redistributive politics. Quarterly Journal of Economics, 121(2), 699–746.
- Berners-Lee, M., Hoolohan, C., Cammack, H., & Hewitt, C. (2012). The relative greenhouse gas impacts of realistic dietary choices. Energy Policy, 43, 184–190.
- Bouchery, E. E., Harwood, H. J., Sacks, J. J., Simon, C. J., & Brewer, R. D. (2011). Economic costs of excessive alcohol consumption in the US, 2006. American Journal of Preventive Medicine, 41(5), 516–524.
- Brehm, J. W. & Cohen, A. R. (1962). Explorations in Cognitive Dissonance. John Wiley & Sons Inc.
- Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. Quarterly Journal of Political Science, 10, 519–578.

- Camerer, C. F. & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. Journal of Risk and Uncertainty, 19(1-3), 7–42.
- Caviola, L., Everett, J. A., & Faber, N. S. (2019). The moral standing of animals: Towards a psychology of speciesism. Journal of Personality and Social Psychology, 116(6), 1011–1029.
- Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. The Journal of Finance, 71(1), 267–302.
- Cook, P. J. & Moore, M. J. (2002). The economics of alcohol abuse and alcohol-control policies. Health Affairs, 21(2), 120–133.
- Cooper, J. (2007). Cognitive dissonance: 50 years of a classic theory. Sage.
- Di Tella, R., Galiant, S., & Schargrodsy, E. (2007). The formation of beliefs: evidence from the allocation of land titles to squatters. Quarterly Journal of Economics, 122(1), 209–241.
- Dickinson, D. L. & Oxoby, R. J. (2011). Cognitive dissonance, pessimism, and behavioral spillover effects. Journal of Economic Psychology, 32(3), 295–306.
- Ehrich, K. R. & Irwin, J. R. (2005). Willful ignorance in the request for product attribute information. Journal of Marketing Research, 42(3), 266–277.
- Elinder, M. (2012). Correcting mistakes: cognitive dissonance and political attitudes in Sweden and the United States. Public Choice, 153(1-2), 235–249.
- Elliot, A. J. & Devine, P. G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. Journal of Personality and Social Psychology, 67(3), 382.
- Elwood, P. C., Pickering, J. E., Givens, D. I., & Gallacher, J. E. (2010). The consumption of milk and dairy foods and the incidence of vascular disease and diabetes: an overview of the evidence. Lipids, 45(10), 925–939.
- Erb, K.-H., Lauk, C., Kastner, T., Mayer, A., Theurl, M. C., & Haberl, H. (2016). Exploring the biophysical option space for feeding the world without deforestation. Nature Communications, 7, 11382.
- Espinosa, R. (2019). L'éléphant dans la pièce. Revue d'Economie Politique.
- Facchini, G., Margalit, Y., & Nakata, H. (2016). Countering public opposition to immigration: The impact of information campaigns.

- Fernbach, P. M., Light, N., Scott, S. E., Inbar, Y., & Rozin, P. (2019). Extreme opponents of genetically modified foods know the least but think they know the most. Nature Human Behaviour, 1.
- Festinger, L. (1962). A theory of cognitive dissonance, volume 2. Stanford University Press.
- Freddi, E. (2017). Do people avoid morally relevant information? Evidence from the refugee crisis.
- Gehrsitz, M. & Ungerer, M. (2017). Jobs, crime, and votes: a short-run evaluation of the refugee crisis in Germany. IZA Discussion Paper, (1094).
- Goetzmann, W. N. & Peles, N. (1997). Cognitive dissonance and mutual fund investors. Journal of Financial Research, 20(2), 145–158.
- Goldsmith, A. H., Sedo, S., Darity Jr, W., & Hamilton, D. (2004). The labor supply consequences of perceptions of employer discrimination during search and on-the-job: Integrating neoclassical theory and cognitive dissonance. Journal of Economic Psychology, 25(1), 15–39.
- Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. Journal of Economic Literature, 55(1), 96–135.
- Graça, J., Calheiros, M. M., & Oliveira, A. (2016). Situating moral disengagement: Motivated reasoning in meat consumption and substitution. Personality and Individual Differences, 90, 353–364.
- Grewenig, E., Lergetporer, P., Werner, K., & Smarzynska Javorcik, B. (2019). Incentives, search engines, and the elicitation of subjective beliefs: evidence from representative online survey experiments.
- Griggs, I. (2018). UK government to run more than 140 campaigns in 2018/19. PR Week, <https://www.prweek.com/article/1463220/uk-government-run-140-campaigns-2018-19>.
- Grigorieff, A., Roth, C., & Ubfal, D. (2018). Does information change attitudes towards immigrants? representative evidence from survey experiments. Representative Evidence from Survey Experiments (March 10, 2018).
- Grossman, Z. & Van Der Weele, J. J. (2017). Self-image and willful ignorance in social decisions. Journal of the European Economic Association, 15(1), 173–217.
- Hestermann, N., Le Yaouanq, Y., & Treich, N. (2018). An economic model of the meat paradox. Working paper.

- Ho, E., Hagmann, D., & Loewenstein, G. (2018). Measuring information preferences. Available at SSRN 3249768.
- Holmes, A. J. & Anderson, K. (2017). Convergence in national alcohol consumption patterns: New global indicators. Journal of Wine Economics, 12(2), 117–148.
- Hu, F. B., Rimm, E. B., Stampfer, M. J., Ascherio, A., Spiegelman, D., & Willett, W. C. (2000). Prospective study of major dietary patterns and risk of coronary heart disease in men. The American Journal of Clinical Nutrition, 72(4), 912–921.
- Joy, M. (2011). Why we love dogs, eat pigs, and wear cows: An introduction to carnism. Conari Press.
- Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. American Economic Review, 90(4), 1072–1091.
- Larsson, S. C. & Orsini, N. (2013). Red meat and processed meat consumption and all-cause mortality: A meta-analysis. American Journal of Epidemiology, 179(3), 282–289.
- Loughnan, S., Bastian, B., & Haslam, N. (2014). The psychology of eating animals. Current Directions in Psychological Science, 23(2), 104–108.
- Mäkelä, K. (1997). Drinking, the majority fallacy, cognitive dissonance and social pressure. Addiction, 92(6), 729–736.
- Matthey, A. & Regner, T. (2011). Do I really want to know? a cognitive dissonance-based explanation of other-regarding behavior. Games, 2(1), 114–135.
- McNally, A. M., Palfai, T. P., & Kahler, C. W. (2005). Motivational interventions for heavy drinking college students: examining the role of discrepancy-related psychological processes. Psychology of Addictive Behaviors, 19(1), 79.
- Micha, R., Michas, G., & Mozaffarian, D. (2012). Unprocessed red and processed meats and risk of coronary artery disease and type 2 diabetes—an updated review of the evidence. Current Atherosclerosis Reports, 14(6), 515–524.
- Morris, J. (2019). 'Count 14' NHS Scotland campaign and debate on the future of alcohol guidelines. Alcohol Policy UK, https://www.alcoholpolicy.net/alcohol_campaigns/.
- Nyborg, K. (2011). I don't want to hear about it: Rational ignorance among duty-oriented consumers. Journal of Economic Behavior & Organization, 79(3), 263–274.
- Orlich, M. J., Singh, P. N., Sabaté, J., Jaceldo-Siegl, K., Fan, J., Knutsen, S., Beeson, W. L., & Fraser, G. E. (2013). Vegetarian dietary patterns and mortality in adventist health study 2. JAMA Internal Medicine, 173(13), 1230–1238.

- Oxoby, R. J. (2004). Cognitive dissonance, status and growth of the underclass. Economic Journal, 114(498), 727–749.
- Pan, A., Sun, Q., Bernstein, A. M., Schulze, M. B., Manson, J. E., Stampfer, M. J., Willett, W. C., & Hu, F. B. (2012). Red meat consumption and mortality: results from 2 prospective cohort studies. Archives of Internal Medicine, 172(7), 555–563.
- Pan, A., Sun, Q., Bernstein, A. M., Schulze, M. B., Manson, J. E., Willett, W. C., & Hu, F. B. (2011). Red meat consumption and risk of type 2 diabetes: 3 cohorts of us adults and an updated meta-analysis. The American Journal of Clinical Nutrition, 94(4), 1088–1096.
- Piazza, J. & Loughnan, S. (2016). When meat gets personal, animals' minds matter less: Motivated use of intelligence information in judgments of moral standing. Social Psychological and Personality Science, 7(8), 867–874.
- Piazza, J., Ruby, M. B., Loughnan, S., Luong, M., Kulik, J., Watkins, H. M., & Seigerman, M. (2015). Rationalizing meat consumption. the 4Ns. Appetite, 91, 114–128.
- Prior, M., Sood, G., Khanna, K., et al. (2015). You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions. Quarterly Journal of Political Science, 10(4), 489–518.
- Quarterly, F. (2016). Quarter 4, october-december 2015. Frontex, Warszawa.
- Rabin, M. (1994). Cognitive dissonance and social change. Journal of Economic Behavior & Organization, 23(2), 177–194.
- Reynolds, J. (2014). Government Ad spend to increase by 22 percent to almost £300m. the Guardian, <https://www.theguardian.com/media/2014/may/14/governemnet-ad-spend-deficit-scottish-referendum-afghanistan>.
- Rothgerber, H. (2014). Efforts to overcome vegetarian-induced dissonance among meat eaters. Appetite, 79, 32–41.
- Schulze, M. B., Hoffmann, K., Manson, J. E., Willett, W. C., Meigs, J. B., Weikert, C., Heidemann, C., Colditz, G. A., & Hu, F. B. (2005). Dietary pattern, inflammation, and incidence of type 2 diabetes in women. The American Journal of Clinical Nutrition, 82(3), 675–684.
- Sekeris, P. & Vasilakis, C. (2016). The Mediterranean refugees crisis and extreme right parties: Evidence from Greece.
- Serra-Garcia, M. & Szech, N. (2019). The (in) elasticity of moral ignorance.

- Sinha, R., Cross, A. J., Graubard, B. I., Leitzmann, M. F., & Schatzkin, A. (2009). Meat intake and mortality: a prospective study of over half a million people. Archives of Internal Medicine, 169(6), 562–571.
- Song, Y., Manson, J. E., Buring, J. E., & Liu, S. (2004). A prospective study of red meat consumption and type 2 diabetes in middle-aged and elderly women: the women’s health study. Diabetes Care, 27(9), 2108–2115.
- Springmann, M., Godfray, H. C. J., Rayner, M., & Scarborough, P. (2016). Analysis and valuation of the health and climate change cobenefits of dietary change. Proceedings of the National Academy of Sciences, 113(15), 4146–4151.
- Steele, C. M., Southwick, L. L., & Critchlow, B. (1981). Dissonance and alcohol: Drinking your troubles away. Journal of Personality and Social Psychology, 41(5), 831.
- Steinmayr, A. (2016). Exposure to refugees and voting for the far-right:(unexpected) results from Austria.
- Tilman, D. & Clark, M. (2014). Global diets link environmental sustainability and human health. Nature, 515(7528), 518.
- WHO (2018). Global status report on alcohol and health 2018.
- Zimmermann, F. (2019). The dynamics of motivated beliefs. American Economic Review.