

An economic model of the meat paradox

Nina Hestermann^{*} Yves Le Yaouanq[†]
Nicolas Treich[‡]

Preliminary and incomplete
December 9, 2017

Abstract

How can individuals care about animals and, at the same time, eat meat? We design a survey study to explore this “meat paradox”. Survey participants ($N = 3054$) underestimate farm animal suffering, and underestimate it more (i.e., are less realistic) when they eat more meat. Building on the literature on cognitive dissonance, we develop a model in which individuals form self-serving beliefs in order to reduce the moral guilt associated with meat consumption. The model characterizes how individuals’ beliefs about animal welfare and their attitude towards information are affected by the economic environment (e.g., price of meat, salience of animal welfare), and by individuals’ preferences (e.g., taste for meat, moral cost of guilt). Several empirical observations are consistent with our model. Keywords: meat price-elasticity, cognitive dissonance, moral dilemma, self-deception, meat paradox, animal welfare

JEL Codes: D72, D81, D83, D84, Z13

^{*}Toulouse School of Economics, nina.hestermann@gmail.com.

[†]Ludwig-Maximilians-Universität, Munich, yves.leyaouanq@econ.lmu.de.

[‡]Toulouse School of Economics, INRA, Université Toulouse Capitole, ntreich@toulouse.inra.fr.

1 Introduction

Meat has been part of the diet of most human societies (Stanford and Bunn, 2001; Burgat, 2017), and is still omnipresent in human culture (e.g., in holiday meals). Yet, the consumption of meat requires animals to be killed and can be viewed as a morally problematic activity. Most people are emotionally disturbed by the animal suffering associated with modern industrial farming practices (Plous, 1993). Individuals who care about the well-being of animals face a trade-off in their consumption choice between the utility derived from eating meat and the psychological disutility implied by their awareness of the suffering inflicted on farm animals.

The objective of this paper is to investigate the relationship between individuals' consumption of meat and their perception of the well-being of animals raised for consumption. We introduce a behavioral model of demand for meat in which individuals form self-serving beliefs in situations of moral dilemma. Consistent with cognitive dissonance theory (Festinger, 1957), these beliefs help reduce the moral guilt associated with their action that has negative consequences on "others", namely on animals. This model enables us to make precise predictions on the link between consumers' preferences and beliefs, as well as on consumers' reaction to exogenous parameters, such as the price of meat, or the salience of animal welfare through (e.g.) media coverage or exposure to videos.

Exploring the determinants of meat consumption is important for several reasons. First, the meat industry represents a significant fraction of our economic activity; in the US, it accounts for about 6% of the US GDP, including direct and indirect revenues (see NAMI, 2013). However, the future of the sector is in question due to the increasing awareness of the negative impacts of meat consumption. Among the different reasons commonly proposed to reduce meat consumption, such as environmental externalities, or negative effects on consumers' health, the moral dimension plays an important role (Ruby, 2012). Increasingly, animal rights organizations provide consumers with information about the harsh living conditions of animals in the farming industry, and animal sciences demonstrate animals' capacity to suffer physically and emotionally (Low et al., 2012; Ritvo, 2007). If innovations such as plant-based proteins or in vitro meat become an afford-

able substitute for meat in the future, the consumers will essentially have to choose between “meat” produced with or without animals. The moral dimension could then be one of the chief determinants of this choice.

Second, meat consumption concerns the daily choices of billions of individuals; it thus offers an important application of a recent literature in psychology and behavioral economics that studies individuals’ behavior in situations of moral dilemmata. This literature, reviewed in Section 2, documents that individuals have a tendency to form self-serving beliefs that understate the negative consequences of their actions on others in order to alleviate the psychological discomfort that results from a selfish action. In the case of meat consumption, this behavioral phenomenon has recently received some attention in psychology under the expression “meat paradox” (Loughnan et al., 2010). We review the relevant literature in psychology in Section 2.

In Section 3, we report the results of a survey designed to elicit individuals’ meat consumption habits and their beliefs about the conditions under which animals are raised. This survey departs from the existing literature in psychology in three ways. First, participants are drawn from a large-scale sample representative of the French population, from which we obtain personal characteristics such as their gender, age and education level. Second, the survey contains questions about the objective reality of the farming industry (e.g., the percentage of pigs raised indoors), which allows us to compare individuals’ beliefs to a true, objective value. Third, and lastly, this survey combines a large set of questions asked separately in previous studies, which allows us to analyze which kind of justifications are appealing to consumers. Our main empirical results are the following. First, there is a negative association between the quantity of meat that individuals consume and the degree to which they believe livestock suffer under industrial farming practices. This result holds both for questions about objective realities and about subjective perceptions. On average, participants tend to underestimate the suffering inflicted on animals compared to objective values. (They underestimate the percentage of pigs raised indoors, for example.)

We explain this empirical association between meat consumption and perception of animal suffering by a theory inspired from the literature on

cognitive dissonance: according to this theory, individuals deny objective information regarding the living conditions of animals in order to alleviate the feeling of guilt that they experience when they consume meat. According to this explanation, consumption of meat causally determines the perception of animal welfare, even among individuals who have received the same pieces of information. Our model works as follows. An individual selects a quantity of meat for consumption at a unit price p . Consuming a quantity c of meat delivers some utility $U(c)$, but inflicts a moral cost ωxc to the agent: ω is the extent to which the individual feels empathetic towards animals and internalizes their suffering, or to which the individual feels guilty about her consumption, whereas x measures the perceived intensity of the suffering inflicted on animals as a result of meat consumption. The variable x is uncertain ex ante, and encompasses all the parameters pertaining to the perception of the consequences of meat consumption: whether animals are able to feel physical and emotional pain, how much livestock suffer from farming practices, the extent to which consuming meat is necessary for good health etc. The core of the model, in line with recent evidence on information processing in situations of moral dilemmata, is that individuals do not care about the “true” value of x but only about their perception of it; as a result, they have incentives to understate their expectation of x in order to keep consuming large quantities of meat without feeling excessively guilty about it.

We model information processing as an intra-personal game, adapting the architecture of the model of motivated beliefs provided by [Bénabou and Tirole \(2002\)](#). The consumer has two selves: Self 0 receives some information about the true value of x , and decides either to transmit the information to Self 1, or to conceal it (at some cost of self-deception, which might vary with the salience of animal welfare). Self 0 thus has an incentive to manipulate the transmission of information to Self 1 in order to distort Self-1’s beliefs about the negative consequences of her actions. We focus on Self 0’s choice on whether to conceal bad news, i.e. signals that indicate a large externality and therefore prescribe a low consumption of meat. In equilibrium, individuals who decide to engage in self-deception are biased in their perception of the suffering of animals, and form different beliefs from those of realistic individuals, even though all agents have received the

same piece of information in the first place.

The model makes several predictions, some of which are supported by the survey study, while others need further and specific investigation. First, there is a negative association between individuals' subjective expectation of x and their consumption of meat, as documented in our survey. This result is due to the fact that individuals with the lower cost of self-deception, or the larger taste for meat, are more likely to engage in self-deception, since the marginal benefit of denying the moral externality is the largest for these individuals; meat eaters are therefore likely to understate the suffering of animals, whereas vegetarians have no incentives to form distorted beliefs about it. In the survey, we indeed find that meat eaters tend to form optimistic beliefs about the conditions in which animals are raised; in contrast, vegetarians' beliefs are closer to the truth.

Second, consumers become more realistic when the price of meat increases; intuitively, an increase in the price of meat makes consumption less appealing, and therefore the incentives to deny the negative consequences of consumption also shrink. The model therefore predicts that variations in prices, an exogenous parameter, influence individuals' perception of the suffering of animals; this contrasts with the standard theory of heterogeneous preferences, in which variations in prices do not have any effect on consumers' beliefs. Observing natural variations of meat prices and their effect on consumers' perception of animal welfare may be a promising way to test the theory with observational data.

Third, we investigate how consumers react to variations in the cost of guilt ω . This comparative statics result is useful for two reasons. First, it allows us to predict which consumers are the most likely to form distorted beliefs. Second, it shows that the effects of advertising campaigns (for instance, by animal activists) meant to increase the social stigma borne by meat eaters depend on the elasticity of the individual's consumption. We show that an increase in ω yields subtle effects, and that the consumer becomes more realistic if and only if her demand for meat is sufficiently elastic. The intuition is that a consumer with an elastic demand can easily adapt to changes in ω by reducing her consumption of meat, which therefore reduces the need to engage in self-deception; by contrast, a consumer whose demand is very inelastic will ultimately decide to maintain the same

consumption level following a variation in ω , and therefore she feels a higher need to alleviate the guilt generated by her consumption. Alternatively, the result means that individuals with a higher propensity to guilt are not necessarily more realistic regarding the effects of their actions; this is true only if their preferences are elastic.

Fourth, we examine the agents' ex-ante attitude towards information. In any standard model, consumers are better off when they receive some information about the consequences of their actions; thus, they should be willing to acquire information about the living conditions of livestock. Things are different in our model, since beliefs directly influence the agent's utility. We show that individuals are information-averse if their equilibrium strategy is to deny bad news. Intuitively, such individuals are harmed by the information, which does not affect their actions but leads them to engage in a more active process of self-deception. This simple observation offers a potential way to test for the basic mechanism of the model, by eliciting whether individuals are willing to receive some objective information regarding the practices of industrial farming or not.

Fifth, and finally, we examine in an extension a model in which $n+1$ identical individuals interact. There are two forms of such interactions. First, individuals observe the meat consumption of others and may learn from that observation the true value of animal suffering. Second, individuals are affected by the aggregate consumption of meat in society: their moral cost depends on the full amount of suffering inflicted on farm animals and not solely on the suffering due to their own consumption. Under such consumption externalities, two symmetric equilibria can coexist, one equilibrium featuring collective realism and one equilibrium of collective denial. There is social reinforcement of denial, in the sense that, under denial, learning opportunities are limited while meat consumption, and thus animal suffering, is high, so that the incentives to stay uninformed are higher. As a result, the existence of individuals who deceive themselves further stimulates denial.

Section 2 reviews the relevant literature. Section 3 exposes the results of our survey. Section 4 introduces the primitives of the model, whose results are presented in Section 5. Section 6 introduces consumption and informational externalities in the model. We conclude the paper in Section

7 by discussing the model and its policy implications.

2 Related literature

Anyone who eats meat is an interested party. For behind the mere momentary desire to eat meat on a particular occasion lie many years of habitual meat-eating who have conditioned our attitudes to animals.

(Peter Singer, *Animal Liberation*, 1975)

In this section we provide a two-part review of the relevant literature; we start with the evidence pertaining the moral aspects of meat consumption and their effects on individuals' beliefs, and we then proceed with the literature on motivated beliefs in moral dilemmata produced by economists and social psychologists.

2.1 Moral ambivalence of meat eaters

A lot of casual evidence suggests that individuals hold empathetic feelings towards animals. For instance, there exist more than 60 million domestic dogs in the US, and many of them are treated as family members, receiving for instance a birthday or a Christmas present (Coate and Knight, 2010). Francione (2004) reports that two-thirds of Americans polled by the Associated Press agree with the following statement: “An animal’s right to live free of suffering should be just as important as a person’s right to live free of suffering”, and more than 50% of Americans believe that it is wrong to make fur coats or to hunt animals. Yet, every year billions of animals are raised in industrial farming facilities for their food, and their fate seems to be largely ignored by the majority of individuals who consume meat on a regular basis. This “moral schizophrenia” (Francione, 2004), or “moral hypocrisy” (Sunstein, 2004), illustrates the apparent discrepancy between individuals’ attitude towards animals and the consequences of their daily actions. It has been the object of a recent strand of academic research in psychology referring to “the meat paradox” (Loughnan et al., 2010). The rest of this section reviews this body of work.

This review starts with studies contrasting the beliefs of vegetarians and omnivores. Omnivores often readily deny that animals used as food may suffer (Bratanova et al., 2011). Rothgerber (2014) stresses that the more respondents endorsed statements such as “animals don’t really suffer when being raised and killed for meat” and “animals do not feel pain the same way humans do”, the more meat they reported consuming. In their survey studies, Bilewicz et al. (2011) report that vegetarians and omnivores attribute different emotional capacities to different animals. They indicate for instance that omnivores, unlike vegetarians, draw a line between the emotions of meat versus non-meat animals (Bilewicz et al., 2011). Meat eaters seem to be poorly informed about the realities of industrial farming. Consistent with singer Paul McCartney’s famous sentence “if slaughterhouses had glass walls, everyone would be vegetarian”, many vegetarians seem to cease consuming meat after a “shock”, for instance after being exposed to the slaughter of an animal, witnessing violence towards animals, or reading influential books such as those of Pollan (2006) and Safran Foer (2010). This suggests that the consumers’ prior expectations about the treatment of animals used for food are biased.

People often justify ill-treatment by denying the ability to suffer. Bastian and Loughnan (2017) put it as follows: “Acting in ways that harm the interests of others is only psychologically problematic to the extent that those others are perceived to have the capacity to suffer. Harm is only aversive when it is directed toward morally relevant others. Smashing a rock, for example, is harmful to the rocks integrity but hardly aversive because rocks are not morally relevant others: they cannot experience harm”. In our context, individuals may change their perception of the animal itself in order to justify the harm done, for instance by belittling the cognitive abilities of animals or their capacity to suffer. The simple act of categorizing an animal as “food” turns out to directly affect how people view that animal (Bratanova et al., 2011). This act may shift people’s focus away from morally relevant attributes (i.e., the capacity to suffer), and as a result change their perception of the animal’s moral worth. Bastian et al. (2012) show that individuals ascribe fewer mental capacities to animals considered appropriate for human consumption. This is illustrated on Figure 1 that displays a negative correlation between the ($N = 63$) subjects’




Figure 1 – Source: [Bastian et al. \(2012\)](#)

perceptions of the edibility and of the cognitive ability of various animals.

Beliefs about animals seem easily malleable, and experimental studies have shown that omnivores’ beliefs are sensitive to trivial psychological manipulations. In one experiment, subjects were asked to eat “beef jerky” or “dried banana” ([Loughnan et al., 2010](#)). They then were invited to a second, supposedly unrelated study about their beliefs about the moral rights of a cow. Subjects who had previously eaten beef reported a lower moral concern than those who had eaten banana. In a follow up series of experiments, [Bastian et al. \(2012\)](#) asked subjects to eat either cold cuts of roast beef or slices of apple. Prior to eating the food, subjects were given the chance to rate the moral capacities of cows. Participants who chose roast beef reported lower moral capacities for cows. Hence, expecting to eat meat in the immediate future is associated with a motivated denial of mind to food animals.

People can use various devices to reduce the feeling of guilt when they consume meat. In an influential work, [Joy \(2011\)](#) presents three categories of justification typically used by meat eaters, known as the “Three Ns justification”. Joy argues that meat eaters come to believe that eating meat is natural, i.e., it is written in our biology; that eating meat is normal, i.e., that it is what most people in civilized society do; and that eating meat

is necessary, i.e., that we need to consume meat in order to be healthy. [Piazza et al. \(2015\)](#) add a fourth N that eating meat is “nice”, and provide experimental evidence for the relevance of these four dimensions for consumers’ choices. In [Te Velde et al. \(2002\)](#)’s survey study, consumers invoke arguments like: “cows and pigs exist to produce meat”, or “a broiler never sees daylight, but it doesn’t know what it misses” or that “it is very important for me to eat meat”. Another strategy, among others, is to build a mental separation between meat and the animal in order to loosen the link between the piece of meat and the dead animal. Indeed, pieces of meat that make this separation difficult (e.g., eyes, tongues, brains, etc.) are often unwillingly handled by consumers ([Kubberød et al., 2002](#)), who are also reluctant to use language that reminds them of the flesh of animals, using neutral substitutes instead (e.g. bacon, hamburger, sirloin). In a recent comprehensive experimental study, [Kunst and Hohle \(2016\)](#) demonstrate that these various dissociation strategies are effective at reducing empathy and at increasing the willingness to eat meat.

2.2 Motivated beliefs in moral dilemma

Scholars from several disciplines (economics, philosophy, psychology) have long debated about the nature of pro-social preferences, with a particular interest in the sincerity of altruism. In recent literature surveys on moral dissonance, [Barkan et al. \(2015\)](#) and [Shalvi et al. \(2015\)](#) discuss how people use pre- or post-justifications to downplay a moral violation. A central question in the literature concerns the devices that people employ to reduce the dissonance, including avoiding information, selecting the sources, misinterpreting the data, or searching for additional evidence ([Von Hippel and Trivers, 2011](#); [Golman et al., 2017](#)). Some experimental evidence indicates that subjects use the uncertainty strategically in order to excuse their opportunistic behavior ([Dana et al., 2007](#); [Exley, 2015](#)), or avoid placing themselves in situations where a pro-social action is expected ([DellaVigna et al., 2012](#); [Andreoni et al., 2017](#)). A phenomenon more directly relevant to our model is that individuals form self-serving beliefs that allow them to rationalize selfish behavior ([Di Tella et al., 2015](#); [Andreoni and Sanchez, 2014](#)), and actively avoid receiving information about the

consequences of their actions on others (Dana et al., 2007; Shu and Gino, 2012; Van der Weele, 2014; Grossman and Van Der Weele, 2017). See Gino et al. (2016) for a survey of this literature.

More generally, our theory builds on a growing literature in behavioral economics that departs from the classical view that beliefs are a function of information only, and that developed models where beliefs are a choice variable that individuals select (under some constraints) in order to satisfy their psychic or instrumental needs. This view that beliefs are “motivated”, i.e. that they include some elements of purposefulness, has been applied to many different questions, such as, among others, risk-taking (Akerlof and Dickens, 1982; Brunnermeier and Parker, 2005), the formation of beliefs about oneself (Bénabou and Tirole, 2002) or about political or scientific facts (Bénabou and Tirole, 2006). See Bénabou and Tirole (2016) for a survey of this literature. A key question in this literature is how to model the self-deception process, i.e. how to specify the objectives and the constraints that people face when they select their beliefs. Our model builds on the architecture proposed by Bénabou and Tirole in their series of papers (Bénabou and Tirole, 2002; Bénabou and Tirole, 2004; Bénabou and Tirole, 2006; Bénabou and Tirole, 2011), in which beliefs are the outcome of an intrapersonal Bayesian equilibrium. But the insights of the theory are robust to other specifications of the process of beliefs formation, such as the one proposed by Brunnermeier and Parker (2005) in which individuals are free to choose their beliefs without any constraint (see the Appendix).

3 Motivating facts: survey evidence

3.1 Methodology

We carried out a large scale survey among French consumers representative of the French adult population. We gathered personal characteristics, such as age, education, and socio-economic status, of about 3000 individuals. Individuals were also asked to answer detailed questions about their habitual consumption of meat and other animal products. We then explored participants’ beliefs on several topics which shape their perception of the suffering inflicted on animals raised for human consumption.


3.2 Overview of the sample

Our sample consists of 3054 adult respondents, representative of the French population, surveyed in July 2017. The data were collected by the private firm BVA, specialists in the collection of consumer surveys and opinion polls. The survey questionnaire, as well as more detailed information on the way the survey was constituted and our sample collected, can be found in the Appendix. There are 1418 male and 1636 female respondents in our sample; the mean age is 45 (with a minimum age of 18 and a maximum age of 90; about 25 percent of participants have an education level equivalent to A-levels; 48 percent went on to obtain a higher education degree. The median household income is 2250 euros (the mean income 2879 euros).


Our main interest lies in investigating the correlation between individuals' consumption behavior and their perception of animal suffering. To this end, we split the sample into three types of consumers: *omnivores*, who report consuming meat at least "several times a week"; *flexitarians*, who report eating meat "a few times a month" or "a few times per year"; and *vegetarians*, who report that they "never eat meat" (this latter group thus also includes vegans). 2645 respondents (86.61 percent) declare eating meat at almost every meal or several times a week, and thus fall into the category of omnivores; 343 (11.23 percent) are flexitarians and 66 (2.16 percent) vegetarians or vegans. An alternative classification of consumer types obtained from our survey is based on the reported duration since the last occasion on which meat was consumed. We display results based on this alternative definition of omnivores, flexitarians and vegetarians in the appendix to underline the robustness of our results.

3.3 Findings


The main finding from this survey is that individuals' beliefs about animal suffering and their consumption of meat are strongly correlated. In this subsection, we present our results on the different types of questions that we asked.




(a) Question 1: *in your opinion, how many egg-laying hens out of 100 are raised in cages?*



(b) Question 2: *in your opinion, how many pigs out of 100 are raised inside barns without access to the outdoors?*



(c) Question 3: *in your opinion, how many rabbits out of 100 die before the day of slaughter?*



(d) Question 4: *in your opinion, how many piglets out of 100 are castrated without anesthesia?*

Figure 2 – Average beliefs by consumer type for the four objective questions; the orange line on each bar covers the 95 percent confidence interval. The red horizontal line indicates the true value.

3.3.1 Objective questions about industrial farming

We first assessed the accuracy of participants’ knowledge about the actual conditions under which farm animals are raised in France today. This is done by asking them to answer some precise quantitative questions (c.f. Figure 2). These questions have the advantage that we can not only evaluate differences in answers between different types of consumers, but also check the degree of accuracy of their answers.

The focus of the four questions reported in Figure 2 concerns a particular aspect of suffering inflicted on livestock, with higher frequencies indicating a higher prevalence of a particular practice or condition painful to the

animals. Higher reported numbers thus indicate a higher estimation of the degree of suffering. On each of the questions which elicit perceived frequencies regarding the harshness of farming conditions, omnivores consistently report significantly lower numbers than vegetarians, with estimations of flexitarians falling in between the two. The order is the same across all questions we ask, and in the majority of cases this difference is highly significant. (The difference between vegetarians and flexitarians is significant at the 0.1 percent level for three out of the four questions; the difference between vegetarians and omnivores is significant at less than 0.1 percent for all of the four questions.)

On top of the four questions reported in Figure 2, we also asked respondents to estimate the number of animals slaughtered for human consumption each day in France, as well as the space that egg-laying hens dispose of in their cages according to the European regulation. While 3 million non-aquatic animals are killed each day in French slaughterhouses, participants vastly underestimate this number: the mean reply is 1 452 722, about half of the true value. The median estimate lies even further from the truth: 50 percent of participants believe that fewer than 100 000 animals are killed in France each day. To help participants estimate the minimal size of cages required by law, we indicated the surface in centimeters squared that is covered by a sheet of paper of format A4 (620 cm^2), and asked them to give their best estimate of the size of space each hen disposes of inside their cage. Their mean estimate is 1624 cm^2 , the median estimate 952 cm^2 , for a true value of 750 cm^2 . Again, participants' beliefs paint a far too rosy picture of the living conditions of livestock.

3.3.2 Justifications for consuming meat

On top of participants' beliefs about the objective realities of modern farming practices, we also asked for their perceptions of farm animals' capacity to experience physical and emotional pain, as this belief crucially shapes the perceived harm inflicted on animals by industrial farming practices. The consumption of these animals may indeed be less of an ethical problem if they are not experiencing their breeding conditions as that painful, because their capacity to experience pain is limited and very differ-

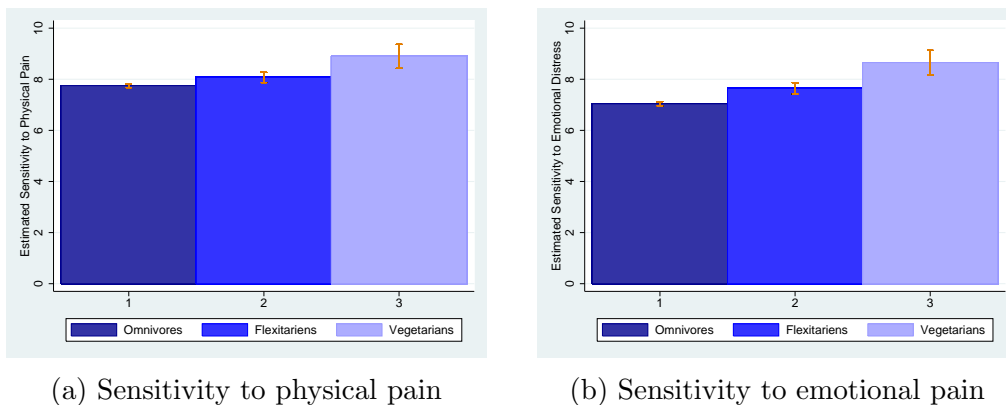


Figure 3 – Average perception of animals’ capacity to experience pain

ent from that of humans for example. Respondents are asked to rate these capacities for different species of animals on a scale from 0 to 10, where 0 represents a capacity for suffering similar to that of inanimate objects, such as rocks, and 10 represents the capacity humans have to suffer physical and emotional pain.

Figure 3 displays respondents’ mean beliefs about the capacity of pigs and hens to suffer pain, for each category of consumers. On both dimensions (physical and emotional), we find that omnivores have a significantly lower perception of this capacity than flexitarians, who in turn score lower than vegetarians. These differences are significant at the one percent level (or lower) between any two consumer groups, for both questions.

In a similar way, if it is perceived as not natural or a deviation from social norms, and thus potentially ostracizing, to not consume meat, this may serve as a justification for not becoming vegetarian. We thus directly ask participants to report, on a scale from 1 to 5, the extent to which they agree with the statements: “eating meat is necessary for staying in good health”, “it is natural to eat meat”, and “eating meat is normal” (5 indicates full agreement, 1 indicates no agreement at all).

Figure 4 displays respondents’ answers for each category of consumers. The graphs show the share of respondents of each type of consumers who either “agree” or “completely agree” (i.e., give an answer larger than 4 on the 5 point-scale) with each of the three statements. As expected, these shares are strikingly different across the three consumer groups, and these differences are again highly significant. (Differences are significant at 0.01

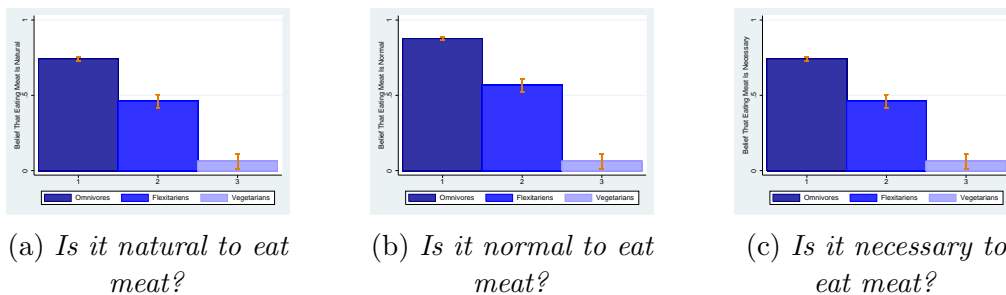


Figure 4 – Average perception of whether eating meat is natural, normal, or necessary.

percent between any two consumer groups.)

3.4 Does consumption justify beliefs, or vice versa?

Our data display a strong association between meat consumption and the perception of animal welfare, and are therefore consistent with a theory in which individuals deny objective information regarding the living conditions of animals in order to alleviate the feeling of guilt that they experience when they consume meat. Note, however, that this association is also predicted by a more traditional Bayesian model in which individuals do not resort to self-deception but consume different amounts of meat because they have received different information about the living conditions of livestock. This theory does not predict that preferences influence perception, but that the causal link is the reverse one, leading from beliefs to consumption decisions.

Here, we present some arguments that speak in favor of the role of cognitive dissonance, which we theoretically explore in Sections 4, 5 and 6. First, the experimental evidence in psychology, reviewed in Section 2, indicates that perceptions of animals (and animal welfare) are affected by variables that should not play any role in a standard model. For instance, Figure 1, quoted from [Bastian et al. \(2012\)](#), shows that people (mistakenly) attribute lower cognitive abilities to edible animals than to non-edible animals, an observation which is difficult to explain by a theory in which beliefs are not manipulated. Moreover, some trivial manipulations (e.g., offering “beef jerky”, or else “dried banana”, as in [Loughnan et al., 2010](#)) affect people’s beliefs about and attitudes towards animals, providing causal evidence that

individuals' perception is indeed affected by their desire to rationalize their consumption.

Second, our data shows a systematic tendency to underestimate animal suffering in the population, as shown by the estimates of the number of animals killed each day, and of the minimal size of cages for egg-laying hens, reported in Section 3.3.1. In addition, our data not only indicate that omnivores, flexitarians, and vegetarians hold different beliefs about the welfare of farm animals, but on the objective questions, vegetarians beliefs are much more accurate on average. As these questions elicit beliefs about verifiable statistics, an objective true value exists for all answers, and we can measure the accuracy of replies for each consumer group. We find that, on average, omnivores significantly underestimate the degree of suffering inflicted on livestock. We calculated an index of the mean of the deviation from the true value across the four objective questions. As reported in Figure 5, this index is about 12 percentage points lower than the true value across the four questions for participants consuming meat frequently. Flexitarians also underestimate this value, but less strongly: their mean difference to the true value is about 9 percentage points, while vegetarians very slightly overestimate the harshness of living conditions along the dimension we elicit: their mean difference lies about 4 percentage points above the true value. The differences between omnivores and flexitarians, and between flexitarians and vegetarians, are significant at 1.4 percent and 0.01 percent, respectively. A large majority of participants of our study underestimate the hardship implied by intensive farming conditions, according to our measure of accuracy of beliefs: 71% of omnivores, 57% of flexitarians, and 38% of vegetarians do so, which together represents 2077 individuals, or about 70% of our sample.

Third, we also observe that the accuracy of beliefs about breeding conditions does not improve with individuals' education level, as displayed for the case of omnivores in Figure 6. Across individuals having obtained one of the four possible education levels (less than high school diploma, high school diploma, 2 year college, and a Bachelor's degree (or more)), beliefs do not vary. Importantly, they do not become more accurate the higher the individual's level of education, as one would expect under a classic model where information acquisition determines consumption choices, and where

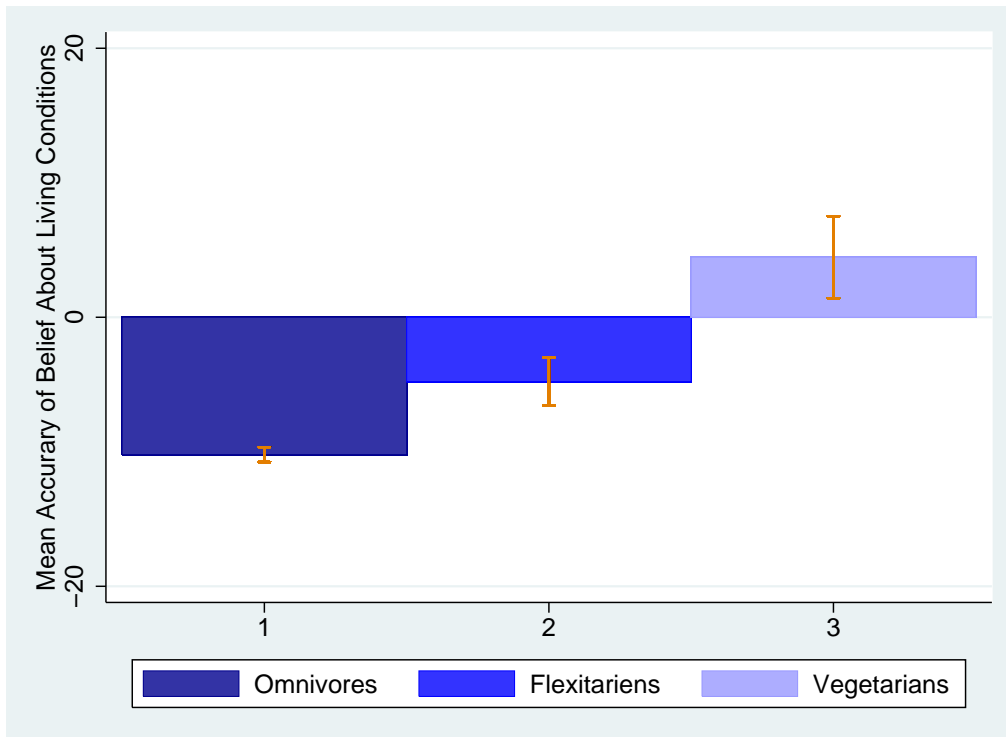


Figure 5 – Mean accuracy of respondents on objective questions

beliefs are not manipulated. More educated individuals can be expected to read more widely and be in general better informed about current debates in the media. They should thus dispose of better knowledge to inform their choices.

These observations speak in favor of a theory in which meat eaters have been avoiding, or repressing, sources of information about the realities of industrial farming, while vegetarians have been gathering and processing the information more accurately.

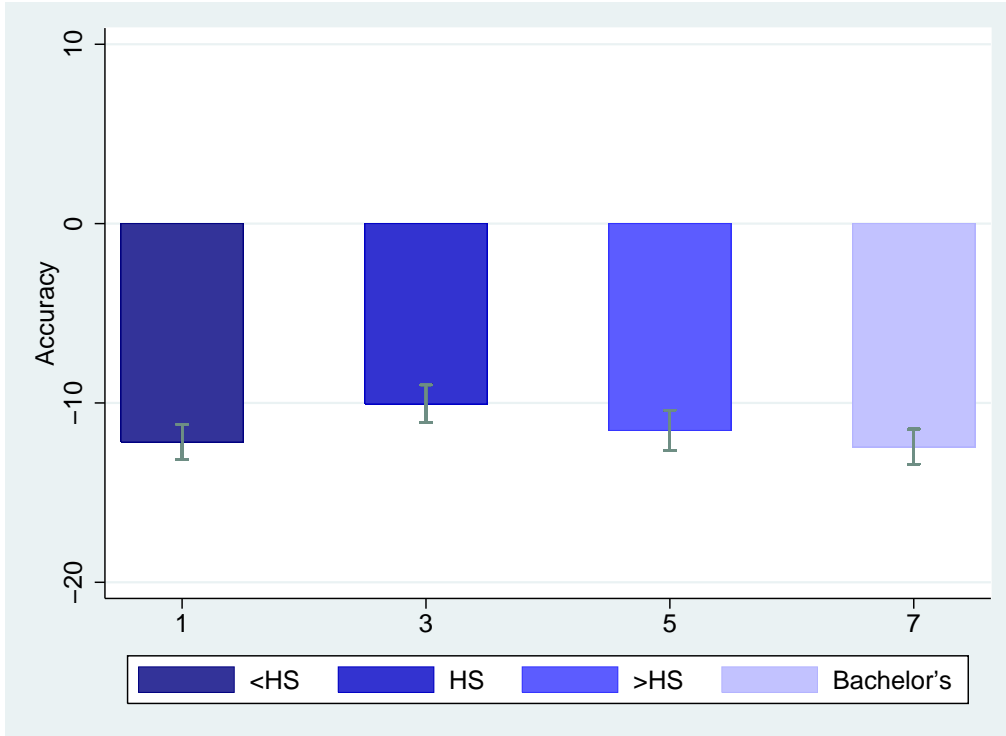


Figure 6 – Mean accuracy of respondents by education level

4 Model

In this section we develop a model based on this theory: consumption of meat causally determines the perception of animal welfare, even among individuals who have received the same pieces of information.

4.1 Primitives

Information and preferences An individual selects a quantity $c \in \mathbb{R}_+$ of meat sold at a price $p \geq 0$. Her taste for meat is represented by the utility function $U : \mathbb{R}_+ \rightarrow \mathbb{R}$ defined over her consumption level c . The function U is twice continuously differentiable and satisfies $U' > 0, U'' < 0$.

On top of the material component $U(c) - pc$, the individual's preferences include some moral concerns: she intrinsically dislikes being responsible for animal suffering. However, she is uncertain about the true level of the externality associated with meat consumption: for instance, she is unsure about the conditions under which animals are raised, about their capacity to feel pain and pleasure, etc. All the uncertainty is captured by a single

random variable X , which takes a high value $X = x_H > 0$ or a low value $X = x_L$ ($0 < x_L < x_H$) with equal prior probabilities. The agent's consumption of meat c generates a negative externality equal to $-Xc$. The states $X = x_H$ and $X = x_L$ convey different moral obligations: the suffering inflicted on animals by meat consumption is greater if $X = x_H$. However, since the individual is uncertain about the state, she evaluates the externality according to her subjective expectation $\tilde{x} = \mathbb{E}X$ and incurs a moral cost of guilt equal to $-\omega\tilde{x}c$. The parameter $\omega \geq 0$ represents the individual's degree of morality, or the extent to which she internalizes animal suffering in her behavior. The individual's prior expectation of X equals $1/2x_L + 1/2x_H$. All in all, c is selected according to

$$\max_{c \in \mathbb{R}_+} U(c) - pc - \omega\tilde{x}c. \quad (1)$$

We write $c^*(\tilde{x})$ for the (unique) optimal solution of this program when the individual has perception \tilde{x} at the time of choice. Equation 1 yields

$$c^*(\tilde{x}) = \max \{(U')^{-1}(p + \omega\tilde{x}), 0\}. \quad (2)$$

The optimal consumption level $c^*(\tilde{x})$ is obviously decreasing with all three variables p, ω and \tilde{x} . We also write

$$V(\tilde{x}) = U(c^*(\tilde{x})) - (p + \omega\tilde{x})c^*(\tilde{x}) \quad (3)$$

for the perceived indirect utility resulting from the individual's consumption decision. Since the externality adds a psychological cost to meat consumption, the function V is nonincreasing in \tilde{x} .

Timeline The formation of beliefs is modeled as resulting from an intra-personal game played by two successive incarnations (Self 0 and Self 1) of the same individual. At date 0, Self 0 receives some information about the externalities of meat consumption. For instance, she might learn about the living conditions of the cattle, watch a video recorded in a slaughterhouse, or learn some information about animals' capacity to experience pain and pleasure. For simplicity, we assume that the signal takes only two values: $m = x_H$ ("bad news") or $m = \emptyset$. The signal is supposed to be perfect, that

is, characterized by $\mathbb{P}[m = x_H \mid X = x_H] = 1$ and $\mathbb{P}[m = x_H \mid X = x_L] = 0$. A signal equal to x_H is thus to be interpreted as an indication that the externalities of meat consumption are large, whereas a signal equal to \emptyset gives support to the alternative hypothesis.

At date 1, once the individual has formed her beliefs \tilde{x} she selects c according to equation 2 and receives her consumption utility. The moral cost is paid at date 2.




Figure 7 – Timeline

Cognitive dissonance In line with the literature on cognitive dissonance, we assume that the individual has the desire and the capacity to distort her perception of the consequences of her actions. When she receives the message $m = x_H$, she might either incorporate this piece of information into her system of beliefs; or, she might incur an effort in order to repress the signal, try to forget about this information, or look actively for contradicting pieces of evidence. We follow the memory-management model proposed by Bénabou and Tirole (Bénabou and Tirole, 2002, 2006; Bénabou, 2013) and model the choice of information processing in an intrapersonal game-theoretic framework. When Self 0 receives the signal $m = x_H$, she can either truthfully transmit $\hat{m} = x_H$ to her future incarnation or repress the signal away and encode $\hat{m} = \emptyset$ at a psychic cost $k \geq 0$ instead. The parameter k (see below, equation (5)) represents the hedonic disutility incurred by the individual if she engages in a rationalization process, denies uncomfortable pieces of evidence, looks for additional information, etc. This cost is likely affected by the degree to which animal welfare is made salient, as discussed in section 1. The more the individual’s awareness is directed towards the topic, the more costly it is for her to suppress information pertaining to it.

The individual is forward-looking when she selects \hat{m} . More precisely, she anticipates that her cognitive choice will affect the consumption level

$c(\tilde{x})$ selected by her incarnation at date 1, as well as the moral cost of guilt $\omega\tilde{x}c(\tilde{x})$ that she will incur. By denying the “bad news” $X = x_H$, the individual alleviates her perceived moral obligations towards animals, and can thus select a higher level of meat consumption.

4.2 Equilibrium concept

The cognitive decision \hat{m} at date 0 is modeled as the equilibrium outcome of the interaction between the individual’s Self 0, who chooses the information, and the individual’s Self 1, who makes the consumption decision. We make two important assumptions regarding this interaction. First, Self 0 fully internalizes Self 1’s perceived utility of consumption without taking into account the true value of the externality. For instance, if Self 0 believes that $X = x_H$ but deceives Self 1 into thinking that $\mathbb{E}X = \tilde{x}(0)$, Self 0 receives the utility

$$U(c^*(\tilde{x}(0))) - (p + \omega\tilde{x}(0))c^*(\tilde{x}(0)).$$

By contrast, if Self 0 were concerned about the true state of the world (and not only about Self 1’s perception of it), she would derive the utility

$$U(c^*(\tilde{x}(0))) - (p + \omega x_H)c^*(\tilde{x}(0)).$$

and there would be no scope for Self-deception: Self 0’s optimal strategy would be to transmit all her information to Self 1 truthfully. Cognitive dissonance arises in this model since the individual is only concerned by her future perceptions when she makes her cognitive decision.

Second, the two players (Self 0 and Self 1) are fully Bayesian. In particular, Self 1 performs Bayesian updating when she does not receive any evidence ($\hat{m} = \emptyset$) and considers the possibility that her Self 0 repressed a signal $m = x_H$. This assumption of sophistication constrains the Self-deception process but does not eliminate the possibility to distort one’s beliefs. Following the message $\hat{m} = x_H$, Self 1’s perception is equal to x_H . Following the message $\hat{m} = \emptyset$ and given Self 0’s habitual strategy σ , Self

1 forms an expectation equal to

$$\tilde{x}(\sigma) := \frac{1-\sigma}{2-\sigma}x_H + \frac{1}{2-\sigma}x_L. \quad (4)$$

The posterior expectation $\tilde{x}(\sigma)$ is decreasing in σ : the higher the individual's habitual probability of transmitting bad news, the more the individual trusts her recollection $\hat{m} = \emptyset$ and believes that the message sent was uninformative. In contrast, an individual who always denies bad news ($\sigma = 0$) stays with her prior beliefs irrespective of her recollection.

We look for Perfect Bayesian Equilibria of this game. An equilibrium is characterized by: (i) Self 0's probability $\sigma \in [0, 1]$ of accepting bad news, i.e. of transmitting $\hat{m} = x_H$ conditional on $m = x_H$; this decision is made optimally conditional on Self 1's behavioral response; (ii) Self 1's consumption decision $c(\sigma, \hat{m})$ at date 1 conditional on the transmitted signal \hat{m} and on Self 0's equilibrium strategy σ . All beliefs are derived according to Bayes' rule. Formally, σ is optimally chosen to satisfy

$$\sigma \in \arg \max_{t \in [0,1]} tV(x_H) + (1-t)V(\tilde{x}(\sigma)) - k(1-t) \quad (5)$$

whereas the consumption level is optimally chosen to satisfy

$$c(\sigma, x_H) = c^*(x_H) \text{ and } c(\sigma, \emptyset) = c^*(\tilde{x}(\sigma)).$$

5 Main results

5.1 Equilibrium behavior

Equation 5 shows that: (i) realism ($\sigma = 1$) is an equilibrium if and only if $V(\tilde{x}(1)) - V(x_H) \leq k$; (ii) denial ($\sigma = 0$) is an equilibrium if and only if $V(\tilde{x}(0)) - V(x_H) \geq k$; (iii) any $\sigma \in (0, 1)$ that satisfies $V(\tilde{x}(\sigma)) - V(x_H) = k$ is an equilibrium strategy.

Equation 4 shows that the benefit $V(\tilde{x}(\sigma)) - V(x_H)$ of wishful denial is nondecreasing with σ , the level of (habitual) realism, strictly so if the individual's consumption $c(\tilde{x}(\sigma))$ is positive. This observation shows the existence and uniqueness of the intra-personal equilibrium.

Proposition 1. *There exists a unique equilibrium of the game characterized by thresholds $0 < k_1 < k_2$ such that: (i) if $k \leq k_1$, the individual denies bad news ($\sigma = 0$); (ii) if $k \geq k_2$, the individual accepts bad news ($\sigma = 1$); (iii) if $k \in (k_1, k_2)$, σ is a strictly increasing function of k .*

In a population with heterogeneous costs of cognitive dissonance, individuals with the lowest k are thus the most likely to resort to willful denial. Consider for instance two individuals, 1 and 2, playing the cognitive strategies $\sigma = 1$ (realism) and $\sigma = 0$ (denial), respectively. Conditional on the message $m = x_H$, individual 1 forms a subjective expectation equal to x_H , and consumes a quantity $c^*(x_H)$, whereas individual 2 forms an expectation equal to $\tilde{x}(0) < x_H$ and consumes a quantity $c^*(\tilde{x}(0)) \geq c^*(x_H)$. In equilibrium, there is thus, as documented in our survey, a negative association between people's subjective expectation of the level of the externality, and their consumption of meat.

In a similar way, for any given individual, under circumstances where the topic is made more salient (e.g. through information displayed on packaging, media coverage, etc.) it becomes harder to deny uncomfortable truths.

5.2 Comparative statics

In this subsection we investigate how the parameters of the decision problem affect the individual's propensity to engage in cognitive dissonance, as characterized by her equilibrium strategy σ , and her resulting consumption decision.

5.2.1 Taste for meat

We start by analyzing the role of individuals' taste for meat in their cognitive choices. Our comparative statics analysis consists in comparing two individuals who only differ in the utility U that they derive from meat consumption. As a first step, suppose that, starting from the belief \tilde{x} and the resulting consumption level $c(\tilde{x})$, the individual can marginally reduce the value of \tilde{x} to $\tilde{x} - \epsilon$. By the envelope theorem, the marginal benefit derived from this change in belief equals $wc(\tilde{x})\epsilon$. This shows that the

benefit from self-deception at any belief \tilde{x} is proportional to the agent's consumption level $c(\tilde{x})$ at this belief. Proposition 2, proved formally in the Appendix, states that a higher taste for meat is associated with a lower level of realism: individuals who (ceteris paribus) consume greater amounts of meat are therefore more likely to resort to willful denial.

Proposition 2. *Individuals with higher demand for meat are more prone to strategic denial: suppose that the utility functions U_A and U_B satisfy $(U'_A)^{-1} \geq (U'_B)^{-1}$, which implies $c_A \geq c_B$; then $\sigma_A \leq \sigma_B$.*

5.2.2 Price increase

In a standard model without self-deception, an increase in the price of meat directly decreases individuals' consumption. In this section, we show that the model of cognitive dissonance also uniquely predicts that an increase in the price increases consumers' awareness of the moral consequences of their actions: the intuition is that a higher price lowers the agents' baseline consumption, making it less profitable to deny bad news in order to rationalize their behavior.

Suppose that, starting from the price p and the beliefs \tilde{x} , and the resulting consumption level $c(p; \tilde{x})$, the unit price of meat decreases marginally to $p - \epsilon$. By the envelope theorem, the marginal utility for the agent resulting from this price increase equals $wc(p; \tilde{x})\epsilon$. The benefit from self-deception, $V(\tilde{x}(0)) - V(x_H)$, is therefore increased by a quantity $w[c(p; \tilde{x}(0)) - c(p; x_H)]\epsilon$, which is nonnegative since $\tilde{x}(0) \leq x_H$: a marginal decrease in the price benefits relatively more to the agent in the scenario where she holds beliefs $\tilde{x}(0)$ than in the scenario where she holds beliefs x_H , as she consumes a higher quantity of meat in the former case. As a result, decreasing the price magnifies the returns to self-deception or, equivalently, the agent's level of realism increases with the unit price of meat. At the limit where $p \rightarrow +\infty$, the agent does not consume any meat irrespective of her beliefs, and therefore she has no incentives to deceive herself ($\sigma = 1$).

Proposition 3. *The individual becomes more realistic when the unit price of meat increases: σ is a nondecreasing function of p . In addition, there exist two thresholds p_1 and p_2 such that $\sigma(p) = 0$ if and only if $p \leq p_1$ and $\sigma(p) = 1$ if and only if $p \geq p_2$.*

In this model, an increase in the price therefore reduces consumption by two channels: a *direct* channel (for fixed beliefs), and an *indirect* channel (through the change in beliefs). The price elasticity of meat consumption is therefore larger than in a model with exogenous beliefs. To see this, consider the thresholds p_1 and p_2 defined in Proposition 3. In the region $[p_1, p_2]$ the individual plays the mixed strategy $\sigma(p)$.¹ A first observation is that, if the price switches from p_1 to p_2 , the individual's perception switches from $\tilde{x}(0)$ to x_H : her consumption response equals $c^*(x_H, p_2) - c^*(\tilde{x}(0), p_1)$, which we rewrite

$$\underbrace{c^*(x_H, p_2) - c^*(x_H, p_1)}_{\text{direct effect}} + \underbrace{c^*(x_H, p_1) - c^*(\tilde{x}(0), p_1)}_{\text{indirect effect}}.$$

Both effects are nonpositive: the indirect effect thus crowds in the standard decrease in consumption that follows from an increase in the price. The total variation is of a larger magnitude than the price response in a model without strategic denial.

Proposition 4. *The aggregate consumption $c(p)$ is nonincreasing with p . In addition, there exist two thresholds $p_1 < p_2$ such that:*

1. *When $p \leq p_1$, $\sigma = 0$ and the individual's perception equals $\tilde{x}(0)$. The price elasticity of her demand is identical to a standard model with unit monetary price $p + \tilde{x}(0)$.*
2. *When $p \geq p_2$, $\sigma = 1$ and the individual's perception equals x_H . The price elasticity of her demand is identical to a standard model with unit monetary price $p + x_H$.*
3. *When $p \in [p_1, p_2]$, σ is an increasing function of p and the individual gradually incorporates moral concerns into her consumption choice. The consumption response $|c(p_2) - c(p_1)|$ is greater than in a model without denial.*

¹Since σ is defined by the implicit equation $V(\tilde{x}(\sigma(p))) - V(x_H) = k$, the implicit function theorem shows that σ is a continuously differentiable function of p in that region.




Figure 8 – Consumption c as a function of the price p

5.2.3 Moral concerns

We now turn to analyzing the effect of an increase in the parameter ω . This parameter can be interpreted as the individual's own feelings of guilt or empathy towards animals, or as the moral pressure imposed by social concerns on the individual's behavior. In the former case, our results speak to whether more altruistic individuals are more likely to form correct beliefs; in the latter case, our results speak to whether aggressive advertising campaigns (e.g., by NGO activists) that aim at increasing moral awareness of animal suffering are likely to reach their objective, that is, reducing consumption.

In this paragraph, we show that this result crucially depends on the elasticity of the individual's demand. More precisely, we show that an increase in the moral cost associated with meat consumption makes individuals more realistic if and only if their consumption of meat is sufficiently elastic. Consumers with elastic demand are those willing to adapt their behavior to a change in the parameters (price and moral cost); when the moral cost associated with meat consumption increases, they react by lowering their consumption, e.g. by becoming vegetarians. In contrast, consumers with inelastic demand are unwilling to reduce their demand; as a conse-

quence, following an increase in the moral pressure, they are more likely to resort to motivated cognition in order to alleviate the cognitive dissonance associated with their behavior.

To see why this result holds, note that the the psychological benefit of playing the strategy $\sigma = 0$ (denying the externality) equals

$$\omega[x_H c^*(x_H) - \tilde{x}(0) c^*(\tilde{x}(0))].$$

Suppose first that $c^*(x_H) = 0$. The individual responds to an increase in her perception of the externality by becoming a vegetarian, which makes her insensitive to any moral pressure. In that case, an increase in ω reduces the benefit from rejecting bad news. Among the consumers who are (hypothetically) willing to become vegetarians if they adopt the signal $m = x_H$, an increase in the moral cost makes the individuals more prone to accept the evidence.

Suppose by contrast that $c^*(x_H) \sim c^*(\tilde{x}(0)) > 0$, meaning that the elasticity of demand is small. In that case, a marginal increase in ω does not trigger any behavioral response but magnifies the psychological cost incurred from meat consumption. Increasing ω therefore gives the individual larger incentives to deny the evidence and avoid thinking about the consequences of her consumption of meat. Among the consumers who have an inelastic demand for meat, those who are the most empathic towards animals are also the most likely to engage in wishful thinking and rationalize the evidence away.

To provide a formal result, let

$$\epsilon_{U'}(c) = -\frac{cU''(c)}{U'(c) - p}$$

be the elasticity of the marginal utility of consumption (net of the unit price) with respect to the consumption level c . This preference parameter is inversely related ($\epsilon_{U'}(c) = 1/\epsilon_{c^*}(x)$ where $c^*(x) = c$) to the elasticity of consumption relative to the perception \tilde{x} , equal to

$$\epsilon_{c^*}(x) = -\frac{\tilde{x}(c^*)'(\tilde{x})}{c^*(\tilde{x})}.$$

Proposition 5. *If $c^*(x_H) = 0$, the individual becomes more realistic (i.e. σ is nondecreasing) when ω increases. If $c^*(x_H) > 0$, two cases arise:*

- (i) *Inelastic demand: if $\epsilon_{U'}(c) > 1$ for all $c > 0$, the individual becomes less realistic (i.e. σ is nonincreasing) when ω increases.*
- (ii) *Elastic demand: if $\epsilon_{U'}(c) < 1$ for all $c > 0$, the individual becomes more realistic (i.e. σ is nondecreasing) when ω increases.*

5.3 Attitude towards information

Finally, we turn to analyzing the individual's attitude towards information from the ex-ante perspective. For simplicity we consider the situation where all parameters of the model are fixed, and we investigate the agent's binary decision of whether to acquire the signal or not.

We show in this section that an individual's attitude towards information depends on her equilibrium cognitive strategy σ : individuals are information-averse if and only if they engage in self-deception. Since incorrect updating is difficult to document empirically, whereas attitudes towards information are more easily elicited, this result offers a possible empirical strategy for detecting motivated reasoning.

For simplicity we focus on the two polar cases of complete realism ($\sigma = 1$) and complete denial ($\sigma = 0$). Suppose first that the individual remains realistic ($\sigma = 1$). Intuitively, there is no conflict of preferences between her two incarnations, who both would like to become better informed in order to make a correct decision. The individual benefits from receiving information, exactly as in a standard model. Formally, if the individual does not receive any information, she stays with her prior perception and her welfare equals $V(\tilde{x}(0))$. If, in contrast, she receives the signal, her posterior beliefs are equal to x_H with probability 1/2 and $\tilde{x}(1)$ with probability 1/2, and her expected welfare thus equals

$$\frac{1}{2}V(x_H) + \frac{1}{2}V(\tilde{x}(1)).$$

Her attitude towards information is therefore pinned down by the convexity of the indirect utility function V in the variable x . Since $V'(x) = -\omega c^*(x)$,

V is convex in x , strictly so in the region where $c^*(x) > 0$. Realistic individuals are therefore information-loving, strictly so if their baseline consumption of meat absent information is positive.

Suppose now that the individual denies bad news ($\sigma = 0$). She forms the same perception equal to $\tilde{x}(0)$, whether she receives some information or not. However, if she receives the information she pays a cost of denial $k > 0$ with probability $1/2$. Intuitively, providing information is useless to the agent since it does not affect her consumption decisions, but it makes her psychologically worse-off due to the cost associated with rationalizing undesirable signals away.

Proposition 6. *Under realism (if $\sigma = 1$), the individual is information-loving. Under denial ($\sigma = 0$), the individual is strictly information-averse.*

6 Extension: consumption externalities

We have so far considered the case of a single agent. To analyze the inter-dependencies between the cognitive strategies and the consumption decisions of different individuals, we now introduce consumption externalities in the model. We assume that individuals' psychological utility is negatively affected not only by their own consumption of meat but also by the aggregate consumption, reflecting the idea that they genuinely care about the welfare of meat animals, irrespective of whether they are responsible for the suffering inflicted.

We assume that the economy is composed of $n + 1$ agents who have identical preference parameters (U, ω, k) and receive a common signal m at date 0, generated as explained in Section 4. The presence of other individuals modifies the analysis of the individual decision problem in two ways. First, we assume common knowledge of rationality. Therefore, at date 2, the individual observes the consumption decisions of her peers, and revises her beliefs about X according to this observation, as the behavior of other individuals might reveal some information about m that the agent has repressed. Importantly, we assume that the agent cannot distort her beliefs about X again if she receives an informative signal, reflecting the idea that the choices of other agents are observed regularly and are more

difficult to rationalize away than the original signal.

Second, the individual is directly harmed by others' consumption of meat. This specification is in line with the interpretation of the model in which individuals have empathic feelings towards animals irrespective of whether the harm is caused by their own or someone else's decisions. Agent i 's utility if she forms an expectation equal to \tilde{x}_i now equals

$$U(c_i) = pc_i - \omega\tilde{x}_i c_i - \xi\tilde{x}_i \sum_{j \neq i} c_j. \quad (6)$$

The parameter ω still represents the cost of guilt caused by the individual's own consumption. The parameter ξ measures the agent's altruism towards animals, and describes the extent to which she internalizes aggregate animal suffering caused by collective decisions.

The timing of the model is as follows. At date 0, all individuals (simultaneously) choose how to process the common signal. At date 1, they form their recollection, update their beliefs, and (simultaneously) make their consumption decisions. At date 2, they observe their peers' consumption decisions, revise their beliefs about X , and receive the utility level given by Equation 6.

For simplicity, we focus on symmetric equilibria in which all individuals play a pure cognitive strategy (either $\sigma = 0$ or $\sigma = 1$), and we assume that the parameters of the model are such that $c^*(\tilde{x}(0)) > 0$, meaning that consumption of meat is positive if individuals stay with their prior beliefs.

Collective realism Let us first study the conditions under which collective realism is an equilibrium. Suppose first that all agents remain realistic and consume $c^*(x_H)$ conditional on the message $m = x_H$. If agent i accepts the evidence $m = x_H$ she receives the utility $V(x_H) - \xi n x_H c^*(x_H)$, where V is defined in Equation 3. In contrast, if she denies the evidence and encodes $m = \emptyset$, she perceives an externality equal to $\tilde{x}(1)$ and thus consumes $c^*(\tilde{x}(1))$. But once she observes the behavior of her peers, who consume $c^*(x_H) < c^*(\tilde{x}(1))$,² she infers that the externality is high and that she has deceived herself into denying the evidence. She then pays a large moral cost for her suboptimally large consumption. Overall, she receives

²Indeed, the assumption $c^*(\tilde{x}(0)) > 0$ implies that $c^*(\tilde{x}(1)) > c^*(\tilde{x}(0)) > c^*(x_H)$.

the utility

$$(c^*(\tilde{x}(1))) - (p + \omega x_H)c^*(\tilde{x}(1)) - \xi n x_H c^*(x_H) < V(x_H) - \xi n x_H c^*(x_H)$$

Realism is therefore a symmetric equilibrium strategy for any $k \geq 0$. In this equilibrium, the incentives to rationalize the evidence away are limited since observing others consuming small quantities makes it impossible for the individual to ignore the externality generated by meat consumption. All in all, the presence of other realistic agents therefore gives the individual higher incentives to accept the evidence.

Collective denial Consider now a candidate equilibrium where all individuals deny the evidence and thus consume $c^*(\tilde{x}(0))$. If agent i deviates and accepts the evidence $m = x_H$, she receives the utility $V(x_H) - \xi n x_H c^*(\tilde{x}(0))$. If she plays the equilibrium strategy $\sigma = 0$, she forms the perception $\tilde{x}(0)$ and observes that all other agents consume $c^*(\tilde{x}(0))$: hence, she does not revise her beliefs about m , knowing that others have deceived themselves as well, and she receives the utility $V(\tilde{x}(0)) - \xi n \tilde{x}(0) c^*(\tilde{x}(0))$. Since all players deny the evidence, their consumption decisions are not informative about the original public message, which increases the scope for self-deception. Denial is an equilibrium strategy if and only if

$$V(\tilde{x}(0)) - V(x_H) + \xi n [x_H - \tilde{x}(0)] c^*(\tilde{x}(0)) \geq k \quad (7)$$

The left-hand side of Equation 7 is increasing in n . Intuitively, the presence of other agents who consume a positive amount of meat $c^*(\tilde{x}(0))$ creates a large psychological cost for any individual; as result, the larger the aggregate consumption $n c^*(\tilde{x}(0))$, the larger the individual incentives to downplay the magnitude of the externality. The existence of deniers thus further encourages denial. As a consequence, since the cognitive strategies of different individuals are strategic complements, multiple equilibria might exist on a non-empty range of the values of the parameters.

Proposition 7. *Suppose that $c^*(\tilde{x}(0)) > 0$. Then collective realism ($\sigma = 1$) is an equilibrium. In addition, there exists $n_0 \in \mathbb{N}$ such that, for all $n \geq n_0$, collective denial ($\sigma = 0$) is also an equilibrium.*

7 Discussion

7.1 Consumers' demand for animal welfare

Experiments or survey studies on consumers' willingness-to-pay for animal welfare identify an interesting paradox (Norwood, 2011): While participants are usually willing to pay a premium for animal welfare that is higher than the cost premium, animal-friendly products only represent a small share on the market. Note that the self-reported concern for animal welfare is not only hypothetical, but survives, to a large extent, to the introduction of real incentives. For instance, Albrecht et al. (2017) measure consumers' concerns for animal welfare directly by eliciting the price that participants are willing to pay to put a living hen into better conditions (organic farming) than intensive farming facilities. While there is substantial heterogeneity in reported willingness to pay, Albrecht et al. (2017) find that 37% of the subjects are willing to pay 25 euros or more, and the average willingness-to-pay equals 14 euros. Participants were thus willing to pay much more than the price premium placed on organically produced poultry in the supermarket in order to save a hen from the suffering it would endure if raised under intensive farming conditions. This difference between consumption choices made when grocery shopping, and the willingness to pay for a hen to live a happier life in an experimental setting is astounding.

The same paradox extends to voting behavior. As an illustration, the referendum concerning Proposition 2 in California in 2008 to reduce the confinement of farm animals³ was accepted by a large majority (63.5% of ballots), in spite of the fact that the conventional products that the law bans (battery cages, veal crates, gestation crates, etc.) largely dominated the Californian market at the time of the referendum. Similar laws were voted in several states in the US. As Norwood and Lusk (2011) put it in their so-called "California egg paradox": "Why are consumers seemingly so compassionate in the voting booth and yet so indifferent in the grocery store?" Our hypothesis to explain these observations is that the salience of

³Proposition 2 stated: "a person shall not tether or confine any covered animal, on a farm, for all or the majority of any day, in a manner that prevents such animal from: a) Lying down, standing up, and fully extending his or her limbs; and b) Turning around freely."

the moral aspects pertaining to the decisions depends on the environment in which individuals make their decisions (grocery store, lab, voting booth). In the grocery store, consumers routinely purchase conventional eggs and meat products, and the living conditions of the animals are left untold by most advertisers. Consumers might therefore divert their own attention away from the moral consequences of their actions at little psychological cost. In contrast, in the experiments mentioned above and in the voting booth, animal welfare is salient, and the consequences of one's actions are unambiguous. As a result, the cost of self-deception increases, and this may lead consumers to adopt a more favorable attitude toward animals. For instance, [Norwood and Lusk \(2011\)](#) document that the demand for free-range eggs has increased by 180% in California during the ballot period. This shows that the ballot made animal welfare salient, and demand for animal welfare then increased, even in the grocery store. ⁴ The meat paradox thus provides an interesting hypothesis to explain the seemingly paradoxical behavior of consumers towards animal welfare.

7.2 Information campaigns

A direct consequence of our model is that various interest groups have large stakes in the transmission of public information about animal welfare. First and foremost, the meat sector has strong incentives to lobby against the provision of information. This seems consistent with some existing laws such as the Animal Enterprise Terrorism Act (AETA) in the US. This law criminalizes the release of images and videos in the farming industry on the grounds of violation of private property. However, the primary motivation for this law may be to hide the reality of farming practices to consumers.⁵ Additionally, the meat sector often provides misleading information about the treatment of farm animals: the packaging of meat products for instance

⁴Voting decisions also differ from individual consumer decisions in that they offer a possibility to solve the coordination problem highlighted in Section 6.

⁵For instance, the New York times ([New York Times, 2013](#)) interprets such “aggag” laws (i.e., antiwhistle blower laws in the agricultural industry) as follows: “The so-called ‘aggag’ laws now being considered by several states, including California, Illinois and Indiana, have nothing to do with protecting property. Their only purpose is to keep consumers in the dark, to make sure we know as little as possible about the grim details of factory farming. These bills are pushed by intensive lobbying from agribusiness corporations and animal production groups.”

typically displays bucolic images of free-ranging and happy farm animals. Our model shows that this supply-side strategy is complemented by the behavior of meat eaters on the demand side, who do not have any incentives to look further into the issue (Proposition 6).

Groups of animal activists have the opposite interests, as they would benefit from informing the population about the objective realities of animal farming. Interestingly, these groups often attempt to directly influence consumers' attitudes through "shock" information campaigns centered on emotionally disturbing content (e.g., comparisons between meat products and human bodies) rather than on statistical information about the meat industry. As the PETA association writes ([People for the Ethical Treatment of Animals \(PETA\), 2017](#)): "It is sometimes necessary to shake people up in order to initiate discussion, debate, questioning of the status quo, and, of course, action." Nevertheless, these tactics seem to have had limited success so far. Our hypothesis is that this strategy, which merely consists in increasing the psychological pressure on meat eaters, might be inefficient (or even backfire) because of the phenomenon uncovered in Proposition 5: increasing the pressure might paradoxically trigger a stronger reaction of denial among the consumers whose preferences are inelastic.

7.3 Humans welfare and animal welfare

We conclude the paper with a short discussion about the choice of the social welfare criterion, and its implications for regulatory practices. An immediate and important question is whether the social planner should be concerned with the "objective" animal welfare or not. Several recent regulatory practices seem to be based on a non-anthropocentric approach ([Bentham, 1789](#); [Singer, 1975](#)) which considers animals as stakeholders in the political decisions.⁶ More subtly, our model features a discrepancy between people's perceptions and the objective reality, and the social welfare criterion must therefore choose whether to take a stand on the accuracy of people's opinions or not. Following the welfarist tradition and the revealed-preference approach, regulators might solely be concerned by the welfare of consumers based on their endogenous beliefs. In that setting, the only

⁶See, e.g., [Johansson-Stenman \(2018\)](#) for some examples.

rationale for policy intervention is the negative externality generated by meat consumption, as consumers are individually selecting their beliefs and actions optimally. In contrast, regulators might be concerned with the “objective” level of animal suffering, and not only by consumers’ endogenous perception. In that case, the preferences of the consumers and of the regulator are not aligned due to the difference in beliefs, so that public intervention would in general be warranted even at the individual level (Salanié and Treich, 2009).⁷

⁷To illustrate these rationales for regulation, consider the California foie gras law (California Senate bill number 1520, approved September 9, 2004). This law prohibits a person from force-feeding a bird for the purpose of enlarging the bird’s liver beyond normal size. In an official commentary to the law, it is said: “According to supporters, the concepts of a free market and free choice assume a fully-informed public, but the vast majority of the public is certainly not fully informed when it comes to the cruelty and suffering involved in the production of foie gras. () The Legislature has already made the sale of the meat of horses and cats and dogs for human consumption a Penal Code violation. It is a tiny step to prohibit the sale of the livers of the birds produced under conditions that cause the birds tremendous suffering”. Hence, the externality due to the suffering of animals as well as the underestimation of this suffering by the population were advanced to justify the regulation.

References

- Akerlof, G. A. and W. T. Dickens (1982). The Economic Consequences of Cognitive Dissonance. *American Economic Review* 72(3), 307–319.
- Albrecht, K., F. Krämer, and N. Szech (2017). Animal Welfare and Human Ethics: A Personality Study. CESifo Working Paper No. 6609.
- Andreoni, J., J. M. Rao, and H. Trachtman (2017). Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving. *Journal of Political Economy* 125(3), 625–653.
- Andreoni, J. and A. Sanchez (2014). Do Beliefs Justify Actions or Do Actions Justify Beliefs? An Experiment on Stated Beliefs, Revealed Beliefs, and Social-Image Manipulation. Working paper.
- Barkan, R., S. Ayal, and D. Ariely (2015). Ethical Dissonance, Justifications, and Moral Behavior. *Current Opinion in Psychology* 6, 157–161.
- Bastian, B. and S. Loughnan (2017). Resolving the Meat Paradox: A Motivational Account of Morally Troublesome Behavior and its Maintenance. *Personality and Social Psychology Review* 21(3), 278–299.
- Bastian, B., S. Loughnan, N. Haslam, and H. R. Radke (2012). Don’t Mind Meat? The Denial of Mind to Animals Used for Human Consumption. *Personality and Social Psychology Bulletin* 38(2), 247–256.
- Bénabou, R. (2013). Groupthink: Collective Delusions in Organizations and Markets. *Review of Economic Studies* 80, 429–462.
- Bénabou, R. and J. Tirole (2002). Self-Confidence and Personal Motivation. *Quarterly Journal of Economics* 117(3), 871–915.
- Bénabou, R. and J. Tirole (2004). Willpower and Personal Rules. *Journal of Political Economy* 112(4), 848–886.
- Bénabou, R. and J. Tirole (2006). Belief in a Just World and Redistributive Politics. *Quarterly Journal of Economics* 121(2), 699–746.
- Bénabou, R. and J. Tirole (2011). Identity, Morals, and Taboos: Beliefs as Assets. *Quarterly Journal of Economics* 126(2), 805–855.

- Bénabou, R. and J. Tirole (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives* 30(3), 141–164.
- Bentham, J. (1789). *An Introduction to the Principles of Morals and Legislation* (Chapters I–V). Wiley Online Library.
- Bilewicz, M., R. Imhoff, and M. Drogosz (2011). The Humanity of What We Eat: Conceptions of Human Uniqueness among Vegetarians and Omnivores. *European Journal of Social Psychology* 41(2), 201–209.
- Bratanova, B., S. Loughnan, and B. Bastian (2011). The Effect of Categorization as Food on the Perceived Moral Standing of Animals. *Appetite* 57(1), 193–196.
- Brunnermeier, M. and J. Parker (2005). Optimal Expectations. *American Economic Review* 95(4), 1092–1118.
- Burgat, F. (2017). *L’Humanité Carnivore*. Editions du Seuil.
- Coate, S. and B. Knight (2010). Pet Overpopulation: An Economic Analysis. *The BE Journal of Economic Analysis & Policy* 10(1 (Advances)), Article 106.
- Dana, J., R. A. Weber, and J. X. Kuang (2007). Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness. *Economic Theory* 33(1), 67–80.
- DellaVigna, S., J. A. List, and U. Malmendier (2012). Testing for Altruism and Social Pressure in Charitable Giving. *Quarterly Journal of Economics* 127(1), 1–56.
- Di Tella, R., R. Perez-Truglia, A. Babino, and M. Sigman (2015). Conveniently Upset: Avoiding Altruism by Distorting Beliefs about Others’ Altruism. *American Economic Review* 105(11), 3416–3442.
- Exley, C. L. (2015). Excusing Selfishness in Charitable Giving: The Role of Risk. *Review of Economic Studies* 83(2), 587–628.

- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.
- Francione, G. L. (2004). Animals—Property or Persons? In C. R. Sunstein and M. C. Nussbaum (Eds.), *Animal rights: Current debates and new directions*. Oxford University Press.
- Gino, F., M. I. Norton, and R. A. Weber (2016). Motivated Bayesians: Feeling Moral While Acting Egoistically. *Journal of Economic Perspectives* 30(3), 189–212.
- Golman, R., D. Hagmann, and G. Loewenstein (2017). Information Avoidance. *Journal of Economic Literature* 55(1), 96–135.
- Grossman, Z. and J. J. Van Der Weele (2017). Self-Image and Willful Ignorance in Social Decisions. *Journal of the European Economic Association* 15(1), 173–217.
- Johansson-Stenman, O. (2018). Animal Welfare and Social Decisions: Is it Time to Take Bentham Seriously? *Ecological Economics* 145, 90–103.
- Joy, M. (2011). *Why We Love Dogs, Eat Pigs, and Wear Cows: An Introduction to Carnism*. Conari Press.
- Kubberød, E., Ø. Ueland, Å. Tronstad, and E. Risvik (2002). Attitudes towards Meat and Meat-Eating among Adolescents in Norway: A Qualitative Study. *Appetite* 38(1), 53–62.
- Kunst, J. R. and S. M. Hohle (2016). Meat Eaters by Dissociation: How we Present, Prepare and Talk about Meat Increases Willingness to Eat Meat by Reducing Empathy and Disgust. *Appetite* 105, 758–774.
- Loughnan, S., N. Haslam, and B. Bastian (2010). The Role of Meat Consumption in the Denial of Moral Status and Mind to Meat Animals. *Appetite* 55(1), 156–159.
- Low, P., J. Panksepp, D. Reiss, D. Edelman, B. Van Swinderen, and C. Koch (2012). The Cambridge Declaration on Consciousness. In Francis Crick Memorial Conference, Cambridge, England, pp. 1–2.

- NAMI (2013). The United States Meat Industry at a Glance. <http://www.meat institute.org/index.php?ht=d/sp/i/47465/pid/47465>. North American Meat Institute. Accessed: 2017-08-15.
- New York Times (2013). Eating with our Eyes Closed. <http://www.nytimes.com/2013/04/10/opinion/eating-with-our-eyes-closed.html>. Accessed: 2017-10-17.
- Norwood, F. B. (2011). The Private Provision of Animal-Friendly Eggs and Pork. *American Journal of Agricultural Economics* 94(2), 509–514.
- Norwood, F. B. and J. L. Lusk (2011). *Compassion, by the Pound: The Economics of Farm Animal Welfare*. Oxford University Press.
- People for the Ethical Treatment of Animals (PETA) (2017). Why does PETA use controversial tactics? <http://www.peta.org/about-peta/faq/why-does-peta-use-controversial-tactics/>. Accessed: 2017-10-17.
- Piazza, J., M. B. Ruby, S. Loughnan, M. Luong, J. Kulik, H. M. Watkins, and M. Seigerman (2015). Rationalizing Meat Consumption. *The 4 Ns. Appetite* 91, 114–128.
- Plous, S. (1993). Psychological Mechanisms in the Human Use of Animals. *Journal of Social Issues* 49(1), 11–52.
- Pollan, M. (2006). *The Omnivore’s Dilemma: A Natural History of Four Meals*. Penguin.
- Ritvo, H. (2007). On the Animal Turn. *Daedalus* 136(4), 118–122.
- Rothgerber, H. (2014). Efforts to Overcome Vegetarian-Induced Dissonance among Meat Eaters. *Appetite* 79, 32–41.
- Ruby, M. B. (2012). Vegetarianism. A Blossoming Field of Study. *Appetite* 58(1), 141–150.
- Safran Foer, J. (2010). *Eating Animals*. Penguin UK.
- Salanié, F. and N. Treich (2009). Regulation in Happyville. *Economic Journal* 119(537), 665–679.

- Shalvi, S., F. Gino, R. Barkan, and S. Ayal (2015). Self-Serving Justifications: Doing Wrong and Feeling Moral. *Current Directions in Psychological Science* 24(2), 125–130.
- Shu, L. L. and F. Gino (2012). Sweeping Dishonesty under the Rug: How Unethical Actions Lead to Forgetting of Moral Rules. *Journal of Personality and Social Psychology* 102(6), 1164.
- Singer, P. (1975). *Animal Liberation: A New Ethics for our Treatment of Animals*. New York: Avon.
- Stanford, C. B. and H. T. Bunn (2001). *Meat-Eating and Human Evolution*. Oxford University Press.
- Sunstein, C. (2004). Introduction: What are Animal Rights? In C. R. Sunstein and M. C. Nussbaum (Eds.), *Animal rights: Current debates and new directions*. Oxford University Press.
- Te Velde, H., N. Aarts, and C. Van Woerkum (2002). Dealing with Ambivalence: Farmers' and Consumers' Perceptions of Animal Welfare in Livestock Breeding. *Journal of Agricultural and Environmental Ethics* 15(2), 203–219.
- Van der Weele, J. J. (2014). *Inconvenient Truths: Determinants of Strategic Ignorance in Moral Dilemmas*. Working paper.
- Von Hippel, W. and R. Trivers (2011). The Evolution and Psychology of Self-Deception. *Behavioral and Brain Sciences* 34(1), 1–16.

Appendix

A Proofs

Proof of proposition 1 Define $k_1 = V(\tilde{x}(0)) - V(x_H)$ and $k_2 = V(\tilde{x}(1)) - V(x_H)$. Both k_1 and k_2 are nonnegative. The function $\sigma \rightarrow V(\tilde{x}(\sigma)) - V(x_H)$ is nondecreasing in σ and maps $[0, 1]$ into $[k_1, k_2]$. Three cases arise:

- (i) If $k_2 \leq k$, $\sigma = 1$ is an equilibrium, uniquely so if $k_2 < k$.
- (ii) If $k_1 \geq k$, $\sigma = 0$ is an equilibrium, uniquely so if $k_1 > k$.
- (iii) If $k \in [k_1, k_2]$, the intermediate value theorem applies since the function $V \circ \tilde{x}$ is continuous, and therefore there exists some $\sigma \in [0, 1]$ such that $V(\tilde{x}(\sigma)) - V(x_H) = k$.

To prove the uniqueness, suppose that $\sigma_1 < \sigma_2$ are both equilibrium solutions. This implies that

$$V(\tilde{x}(\sigma_1)) \geq V(\tilde{x}(\sigma_2)). \quad (8)$$

Indeed, $\sigma_2 > 0$ implies that $V(\tilde{x}(\sigma_2)) - V(x_H) \leq k$, whereas $\sigma_1 < 1$ implies that $V(\tilde{x}(\sigma_1)) - V(x_H) \geq k$.

Since $V \circ \tilde{x}$ is strictly increasing on the range of σ that delivers positive consumption levels $c^*(\tilde{x}(\sigma))$, equation 8 is possible only if $c^*(\tilde{x}(\sigma_1)) = c^*(\tilde{x}(\sigma_2)) = 0$. As a consequence, the monotonicity of c^* implies that $c^*(x_H) = 0$, and therefore $V(\tilde{x}(\sigma_i)) - V(x_H) = 0 < k$ for $i = 1, 2$. This contradicts $V(\tilde{x}(\sigma_1)) - V(x_H) \geq k$.

Proof of proposition 2 Let $c_A(\tilde{x}) = \max[(U'_A)^{-1}(p + w\tilde{x}), 0]$ and c_B be defined similarly. Let σ_A be agent A 's equilibrium strategy. At this equilibrium level, agent A 's benefit from self-deception equals $V_A(\tilde{x}(\sigma_A)) - V_A(x_H)$. The crucial step in the proof is to show that

$$V_B(\tilde{x}(\sigma_A)) - V_B(x_H) \leq V_A(\tilde{x}(\sigma_A)) - V_A(x_H). \quad (9)$$

Once this is established, the result $\sigma_B \geq \sigma_A$ follows immediately from the analysis of the intra-personal equilibrium in proposition 1.

Note that for inequality 9 to hold, a sufficient condition is that the function $x \rightarrow V_B(\tilde{x}) - V_A(\tilde{x})$ is nondecreasing in \tilde{x} , since $\tilde{x}(\sigma_A) \leq x_H$. But, by the envelope theorem,

$$\frac{\partial V_A(\tilde{x})}{\partial \tilde{x}} = -\omega c_A(\tilde{x})$$

and, similarly,

$$\frac{\partial V_B(\tilde{x})}{\partial \tilde{x}} = -\omega c_B(\tilde{x}).$$

As a consequence,

$$\frac{\partial [V_B(\tilde{x}) - V_A(\tilde{x})]}{\partial \tilde{x}} = \omega [c_A(\tilde{x}) - c_B(\tilde{x})] \geq 0.$$

This proves the result.

Proof of proposition 3 Consider a price level p and the resulting equilibrium cognitive strategy σ . At this equilibrium, the agent's benefit from self-deception equals $V(\tilde{x}(\sigma); p) - V(x_H; p)$. By the envelope theorem,

$$\frac{\partial [V(\tilde{x}(\sigma); p) - V(x_H; p)]}{\partial p} = c(x_H) - c(\tilde{x}(\sigma)) \leq 0$$

and thus the net benefit from self-deception is nonincreasing with the price: this observation, together with the analysis of the intra-personal equilibrium in Proposition 1, proves the result.

Proof of proposition 5 Applying the envelop theorem yields

$$\frac{\partial [V(\tilde{x}(\sigma)) - V(x_H)]}{\partial \omega} = -\tilde{x}(\sigma)c^*(\tilde{x}(\sigma)) + x_H c^*(x_H). \quad (10)$$

If $c^*(x_H) = 0$, this expression is nonpositive, therefore σ is nondecreasing in ω .

Suppose otherwise that $c^*(x_H) > 0$, which implies that for all $\tilde{x} \leq x_H$, $c^*(\tilde{x}) = (U')^{-1}(p + \omega \tilde{x})$. If the function $\tilde{x} \rightarrow \tilde{x}c^*(\tilde{x})$ defined on $[x_L, x_H]$ is strictly increasing, expression 10 is nonnegative and therefore σ is nonin-

creasing in ω . This is true if

$$-\frac{\tilde{x}(c^*)'(\tilde{x})}{c^*(\tilde{x})} < 1 \Leftrightarrow -\frac{\omega\tilde{x}}{U''[(U')^{-1}(p + \omega\tilde{x})](U')^{-1}(p + \omega\tilde{x})} < 1 \quad (11)$$

$$\Leftrightarrow -\frac{U''[(U')^{-1}(p + \omega\tilde{x})](U')^{-1}(p + \omega\tilde{x})}{\omega\tilde{x}} > 1 \quad (12)$$

$$\Leftrightarrow \epsilon_{U'}[(U')^{-1}(p + \omega\tilde{x})] > 1. \quad (13)$$

The case $\epsilon_{U'}(c) < 1$ is symmetric. This proves the result.

B An alternative model of self-deception

In this appendix, we use the cognitive dissonance model proposed by Brunnermeier and Parker (2005). Although the model is significantly different from the model presented in the main text, we will show that the key parameters can be interpreted in a similar fashion, and that the main predictions remain analogous.

As above, let the indirect utility defined as follows

$$\begin{aligned} V(x) &= \max_{c \geq 0} u(c) - pc - wxc \\ &= u(c(x)) - pc(x) - wxc(x) \end{aligned}$$

Let $x(t) = tx_h + (1-t)x_l$ the subjective expectations about animal suffering. Suppose now that the consumer observes that the level of animal suffering is x_h . Following Brunnermeier and Parker (2005), the beliefs' choice problem can now be written as follows

$$\max_{t \in [0,1]} (1-k)V(x(t)) + k[u(c(x(t))) - pc(x(t)) - wx_h c(x(t))] \quad (14)$$

This model describes a “compromise” between maximizing the subjective welfare $V(\cdot)$ and maximizing the “objective” welfare as represented by the term into brackets $[u(c(x(t))) - pc(x(t)) - wx_h c(x(t))]$. This last term is coined objective welfare because it accounts for the observed level of animal suffering x_h and is maximized for beliefs $t = 1$. Hence by choosing $t \leq 1$ the agent incurs a welfare loss. In other words, the cost of denying reality is the reduction in objective welfare associated with a suboptimal consumption decision (i.e., too much meat consumption $c(x(t)) \geq c(x_h)$).

The parameter $k \in [0, 1]$ is the weight given to the objective welfare. Note that model (14) can also be rewritten as

$$\max_{t \in [0,1]} V(x(t)) - kwc(x(t))(x_h - x_l)[1 - t] \quad (15)$$

Hence, the choice of optimal beliefs may be viewed as a trade off between the subjective welfare $V(x)$ and a cost of self-deception captured by the term $kwc(x(t))(x_h - x_l)[1 - t]$. Interestingly, this last term increases with the extent of self-deception $1 - t$. Note that this term also increases with

consumption $c(x(t))$ and the parameter k . Hence, the parameter k can still be interpreted as the (per consumption) unit cost of self-deception as in the model considered in the main text.

Differentiating (15) with respect to t we get

$$-(1-k)wc(x(t))(x_h - x_l) - kwc'(x(t))(x_h - x_l)^2[1-t] \quad (16)$$


Note that the first term is negative while the second term is positive under $t \leq 1$ given $c'(\cdot) \leq 0$. In particular, when k is large enough then the second term always dominates, and the maximum must be $t = 1$. Conversely, when k is low enough then the first term always dominates and the maximum must be $t = 0$. Moreover, when an interior solution exists, then the optimal t increases in k since (16) increases in k . As a result, we retrieve a similar characterization of optimal beliefs as in Proposition 1. Moreover, the relationship between optimal consumption and the unit cost of self deception has also a similar form as the one depicted on the Figure 2.⁸ In words, less meat consumption is associated with more realism about animal suffering.

The negative association between consumption and beliefs documented in the survey is thus also predicted by this alternative model.


⁸Take for instance $p = w = x_h = 1$, $x_l = 0$ and $u(c) = \log c$. Then optimal beliefs are equal to $t = \max(0, 2k - 1)$ so that optimal consumption is equal to $c(x(t)) = \min(1, 1/2k)$ for $k \in [0, 1]$.

C Robustness check: alternative definitions of consumer types


In this appendix we reproduce the main empirical results of the survey with an alternative definition of the consumer types, based on the (self-reported) last date of meat consumption before the survey. Omnivores are now defined as those participants who state having consumed meat “today or yesterday” or “a few days ago”. Flexitarians are those who last ate meat “some weeks ago” or “some months ago”, and vegetarians those who declare that “it has been years since I last ate meat”. Applying these definitions, 94.27 % of respondents are classified as omnivores, a higher percentage than under the alternative classification (86.61%). 3.73% are defined as flexitarians, and 2% as vegetarians. Again, the resulting graphs are very similar to those of section 3, with differences between omnivores and flexitarians being slightly less pronounced.



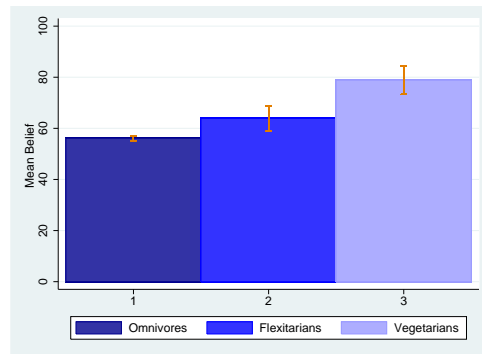
(a) Question 1: *in your opinion, how many egg-laying hens out of 100 are raised in cages?*



(b) Question 2: *in your opinion, how many pigs out of 100 are raised inside barns without access to the outdoors?*




(c) Question 3: *in your opinion, how many rabbits out of 100 die before the day of slaughter?*




(d) Question 4: *in your opinion, how many piglets out of 100 are castrated without anesthesia?*

Figure 9 – Average beliefs by consumer type for the four objective questions; the orange line on each bar covers the 95 percent confidence interval. The red horizontal line indicates the true value.



(a) Sensitivity to physical pain



(b) Sensitivity to emotional pain

Figure 10 – Average perception of animals' capacity to experience pain

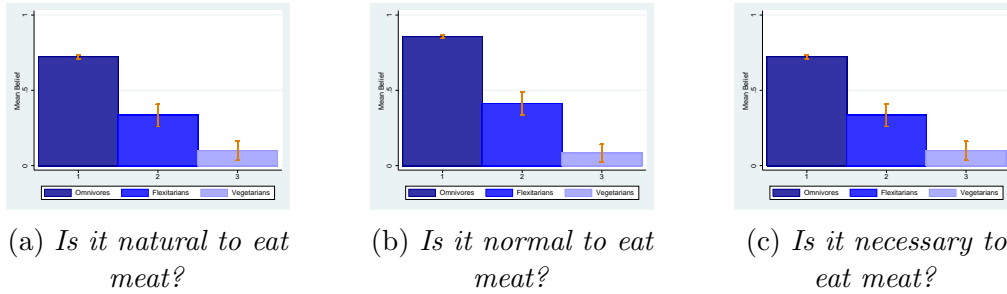


Figure 11 – Average perception of whether eating meat is natural, normal, or necessary.





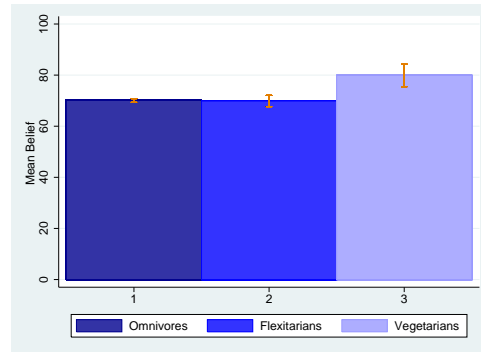
Figure 12 – Mean accuracy of respondents on objective questions

D Robustness check: reduced sample without outliers


In this appendix we reproduce the main empirical results of the survey with a sample where we excluded participants who gave unrealistic answers, suggesting either a lack of attention or a misunderstanding of the questions. The following exclusion criteria are applied: the observations are dropped if (i) the participant believes that 100 or fewer animals are killed each day; (ii) the participant believes that the surface of the cages of egg-laying hens is equal to 10cm^2 or less; (iii) the participant believes that 90% or more rabbits die before the day of slaughter; or, (iv) the participant gave the same reply to all four objective questions. This leaves us with 2535 observations. The graphs listed below are essential similar to the ones in section 3. The ordering of the three consumer groups is identical, and differences between consumer groups are of similar magnitude and significance. We note a slight difference in the mean accuracy of beliefs about livestock suffering, as measured by the mean reply to the four objective questions: vegetarians now give an estimate that is correct on average, i.e. not significantly different from zero.




(a) Question 1: *in your opinion, how many egg-laying hens out of 100 are raised in cages?*



(b) Question 2: *in your opinion, how many pigs out of 100 are raised inside barns without access to the outdoors?*




(c) Question 3: *in your opinion, how many rabbits out of 100 die before the day of slaughter?*




(d) Question 4: *in your opinion, how many piglets out of 100 are castrated without anesthesia?*

Figure 13 – Average beliefs by consumer type for the four objective questions; the orange line on each bar covers the 95 percent confidence interval. The red horizontal line indicates the true value.



(a) Sensitivity to physical pain



(b) Sensitivity to emotional pain

Figure 14 – Average perception of animals' capacity to experience pain

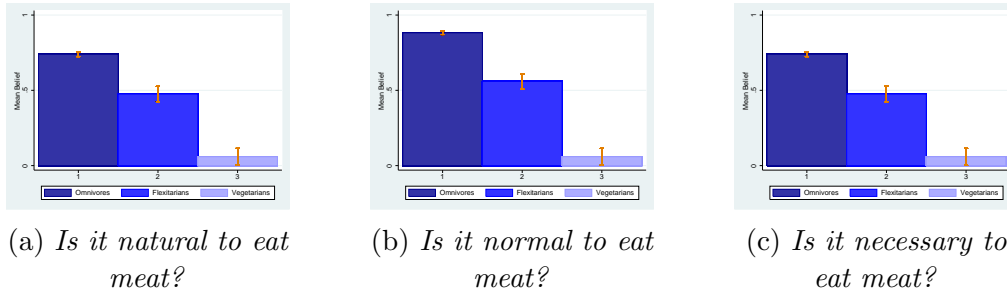


Figure 15 – Average perception of whether eating meat is natural, normal, or necessary.

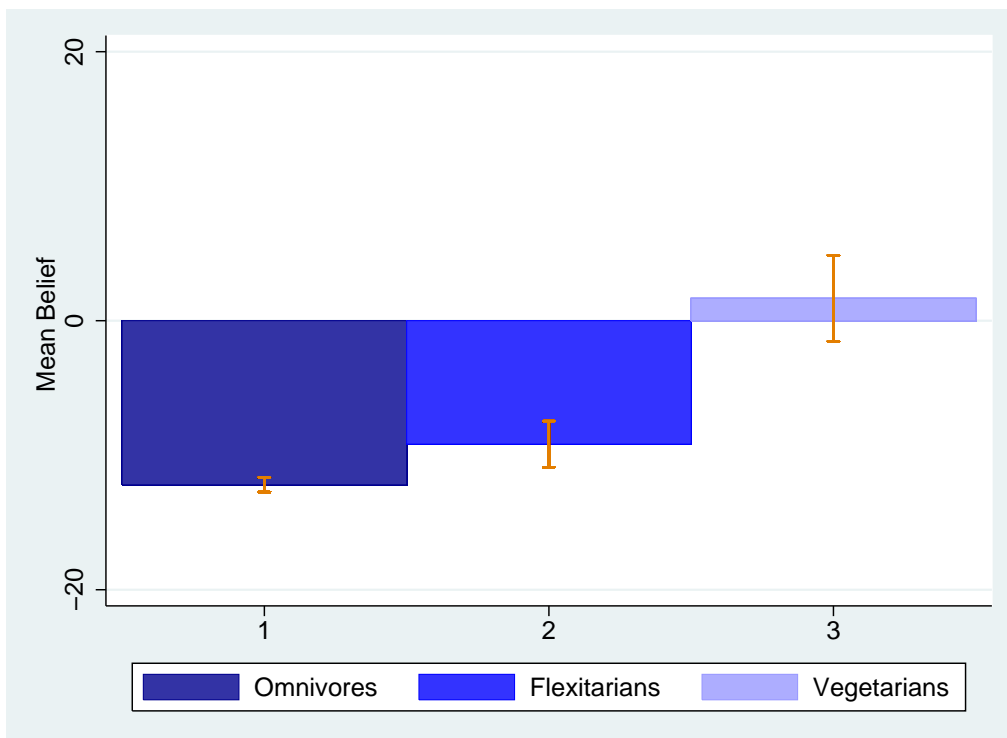


Figure 16 – Mean accuracy of respondents on objective questions